



Innovative Technology
In System Packaging

the Enclosure

One Thing Separates a Great System Enclosure from a Good One.

It's the innovative application of cutting-edge technology to your specific requirements. It's what we do at Electronic Solutions.

If you're the engineer responsible for the system packaging, you have your hands full dealing with power distribution, cooling, and safety concerns such as UL, CSA and VDE approvals. That's where Electronic Solutions can make your job easier! You can rely on our experience and expertise to quickly resolve the design issues you're facing—at a price that even your bean counters will be pleased with.

Airflow testing. Our *in-house* airflow/temperature testing facility

allows us to carefully examine the environment *inside* a system enclosure, testing the effects of wire harness placement, fan size and location, and the characteristics of the airflow path itself. We can even test your complete system to fine tune air direction, cool hot spots, and reduce audible fan noise and power consumption.

Safety Agency Approvals.

Bypass the costly maze of paperwork and testing! We can provide enclosures that are already UL listed, CSA certified, and GS approved through TUV certification. That big step allows you to install recognized boards and peripherals and win

agency approval of your finished system with no additional effort or cost! As a self-testing facility for the agencies, we're equipped to handle every aspect of certification to reduce your costs and shorten your time to market.

EMI/RFI Testing. We subject our enclosures to *in-house* testing per MIL-STD-285, providing an attenuation profile of each unit. When it comes time for you to pursue FCC or VDE approval of your system, we can perform these tests for you and resolve any problems immediately.

Whether it's off-the-shelf or something special that's off-the-drawing board, there's only one way to get the enclosure that's just right for your system. That's to call Electronic Solutions and put our experience to work for you.



VMEbus
VXI bus
Sun bus
Mixed buses
Proprietary buses
Multibus I
Multibus II
Futurebus+

Tabletop,
Rackmount,
DeskMate,
Bellcore and
Custom
Designs



6790 Flanders Drive, San Diego, CA 92121
(619) 452-9333 FAX: 619-452-9464
Call Toll Free: (800) 854-7086
In Calif: (800) 772-7086



*"Synergy's
dual '040
outperformed all
other boards
I've evaluated in
the last five
years."*

—KOR Electronics
Garden Grove, CA

You'll find Synergy products built into the world's most demanding applications.

The Goal:


Design the world's only multi-frequency Radar Target Generator System able to simulate hostile threats on the military's diverse radar systems.

The Problem:

To get the needed horsepower, sixteen 68040 CPU boards were required. However, with this many boards, VMEbus bandwidth limits would severely degrade system performance making the project unfeasible.

The Solution:

Synergy's V420 dual 68040 SBC. After evaluating several products, KOR Electronics selected eight V420s which could deliver 320 MIPS without VMEbus bandwidth degradation.



*"Not only did
we meet our project
performance goals,
we were able to
reduce our costs
by 40%."*

Unexpected Benefits:

Synergy's dual '040 solution cut KOR's hardware requirements by 50% while vastly increasing system reliability. These unexpected benefits reduced system costs by 40%.

In Synergy, KOR also found a design partner with strong integration expertise and dependable customer support.

"I recommend you call Synergy today."

Next time you need high performance SBCs, do as KOR Electronics did. Call Synergy Microsystems. You'll be as satisfied as they are.




High performance SBCs for demanding applications.

SYNERGY MICROSYSTEMS, Inc. 179 Calle Magdalena, Encinitas, CA 92024, 619-753-2191, Fax 619-753-0903

DESIGN *and Design Directions*

FOR ELECTRONIC ENGINEERS & ENGINEERING MANAGERS



In the days of the '49ers, it was gold that was in "them thar hills." Today, up the road from Sutter's Mill, it's silicon—and the latest rush is on—to analog design.107

Illustration by
Bill Morrison

Editorial.....20

Advertiser Index.....150

COMING NEXT MONTH

Designing for testability

Mezzanine bus strategies

Advances in datacom ICs

Emulators

Robotics and automation

NEWS BRIEFS

ovi and vi get down to business . . . Intel deemphasizes Multi-bus . . . Altera supports user-selected design tools . . . Bus interface chips hit the street . . . Next finds big buy for 486 operating system . . . Pact eases setup of X Windows on PCs to networks . . . Intel enters FPGA game . . . FPGA makers minimize handling.....10

TECHNOLOGY DIRECTIONS

Computers & Subsystems

First Futurebus+ to VME bridge board appears.....30
MOSAS wants military to drive commercial standards 38

Software & Development Tools

Real-time environments aimed at cost-sensitive 32-bit designs44

Integrated Circuits

Chip sets could unleash workstation clones.....48
Specialty PROMs count bursts, implement state machines52
Leaders in mixed-signal design to meet in San Francisco.....57

NEW PRODUCT DEVELOPMENTS

ASICs & ASIC Design Tools

Register-rich FPGA family has 9-ns interconnect delays.....134

Computers & Subsystems

DSP chip set aimed at real time.....135
IPs get PCMCIA adapter137

CAE/CAD Design Tools

EDA suite boasts front-to-back tool integration138

Software & Development Tools


UNIX-style communications for embedded systems140

SPECIAL REPORT ON FUTURE COMPUTING


Is neural computing the key to artificial intelligence?

Neural networks are touted as a technology that can predict everything from the outcome of a horse race to loan eligibility. Are we witnessing an artificial intelligence revolution or only headline-grabbing hype?

— Mike Donlin and Jeffrey Child 87



Page 30



Page 48



Page 87

SOURCEGATE

integrating high level language debugging with in-circuit emulators.

SourceGate is a window driven high level language debugger designed to support the Huntsville Microsystems 200 series of in-circuit emulators.

- User configurable windows can be sized, moved and duplicated anywhere on the screen.
- Code can be viewed in all displays (trace, single step, etc.) in one of three modes: Source only, Assembly only or both Source and Assembly.
- Watch windows display and monitor code variables.
- Optional Performance Analysis Card for real-time software performance analysis and real-time software test coverage.
- Available for IBM PC family and UNIX systems including Apollo and SUN.

For more complete technical information, write to Huntsville Microsystems Inc., 3322 South Memorial Parkway, Huntsville, AL 35801 or call (205) 881-6005. Ask for *free* demo disk.


IBM is reg. T.M.
International Business Machines, Inc.
Unix is reg. T.M.,
Bell Laboratories, Inc.



Available Emulators

| | |
|------------|-------------------|
| 68000 | 68EC020 |
| 68008 | 68EC030 |
| 68010 | 68HC001 |
| 68020 | 68HC11 Family |
| 68030 | including F1&D3 |
| 68302 | 8051 Family |
| 68301/303 | DS5000 |
| 68330/333 | 8096/80196 Family |
| 68331/332 | 8085 |
| 68340 | 64180/Z180 |
| 6809/6809E | Z80 |
| | 68HC16 Family |

Now supporting 68040 Series



Tiburon's Greg Pope on: Software testphobia

While the United States is the unequivocal leader in the global software market, it's no secret that our software is infested with bugs. 25

Look for the premier of
OEM
 INTEGRATION
next month, featuring a close-up analysis of the microdrive vs. PCMCIA competition

TECHNOLOGY & DESIGN REPORTS

Performance analysis spots hardware/software bottlenecks

If your system works, but doesn't perform as expected, try performance analysis tools. They can provide a high-level view of system operation, zoom in on problems and help you avoid performance bottlenecks in the design stage.

—Tom Williams 61

Windows demands drive PC chip set features

PC chip set makers are focusing on Windows graphics issues. How well they handle acceleration, cache and power management will decide whether or not they prosper in the OEM market. —Don Tuite..... 73

COVER STORY

Analog designers still trail behind their digital counterparts

Still dealing with custom more than semicustom, analog designers plod along the trail without the help of synthesis, a standard analog HDL or test methodology. But the determined can still hit pay dirt with the tools at hand.

—Barbara Tuck107

DESIGN STRATEGIES

Xerox brings cost-efficient strategy to The Reading Edge

Xerox Imaging Systems redesigns The Reading Edge, a machine that converts printed characters into synthesized speech for the blind and visually impaired. —Jeffrey Child123

PRODUCT FOCUS

High-speed A-D converters shift to new architectures

Driven by the demands of applications such as professional video, radar, high-definition television, and ultrasound, designers have an endless appetite for faster analog-to-digital conversion. —Jeffrey Child129


COLUMN

MIXED-SIGNAL DESIGN — Stephan Ohr


AHDL moves slowly toward standardhood142



Page 61



Page 73




Page 129

Put These G And Take So



Think of it as peer pressure at its finest. All these companies want you to take some time off: Time off your design cycle, and your time to market. They're the Fusion29KSM Partners, and their support helps make the 29KTM Family of embedded RISC processors so fast and easy to use.

Fusion29K Partners include over 100 reputable companies with more than 200 different support products. Together,



The 29K Family includes parts for every conceivable application—from the brand new low cost, easy-to-use Am29200TM to the top-of-the-line Am29050TM with on-chip floating point coprocessor.

they provide everything you need to speed up the design process, every step of the way. From evaluation, to system and hardware design, software development, hardware and software integration, system debug—all the way to production.

The Fusion29K Partners have a whole arsenal of time and labor saving devices. You'll find emulators, simu-

ways TO WORK me Time Off.




5M Consulting

AMP

TeleSoft International, Inc.



MICRO-PROCESSOR SERVICES, INC.



TERADYNE

Schneider & Koch



Research



Graphics Resource Group

FLUKE®



cygnus



Epsilon Industries


telenetworks
SOFTWARE FOR DATA AND TELECOMMUNICATIONS A LHM COMPANY



Capricorn Systems
International



HDE, INC.



PEERLESS



Gibbons & Associates

Augat

Custom Applications Inc.

GSI Techsi



EDI

LOGIC
MODELING

FUJITSU



Information &
Communications Mgmt.



XDI

Tektronix
COMMITTED TO EXCELLENCE

IBM



lators, CASE and CAE tools, C compilers, machine and source level debuggers, logic analyzers, and manufacturing support of all kinds. All of the highest quality and performance standards.

And it won't take a lot of time to get acquainted with these development tools, since they include most of the hardware and software products you already know.

For more information on the 29K Family of embedded RISC processors and Fusion29K support, call AMD today: **1-800-292-9263 Ext. 3.**


With this much support, you may be able to take some time off for yourself.



Advanced Micro Devices

901 Thompson Place, P.O. Box 3453, Sunnyvale, CA 94088. © 1992 Advanced Micro Devices, Inc.
29K is a trademark and Fusion29K is a servicemark of Advanced Micro Devices. All other brand or product names are trademarks or registered trademarks of their respective holders.

This is your chance to get the full power of embedded



FORCE Computers, Inc., 3165 Winchester Blvd., Campbell, CA 95008-6557, Prof.-Messerschmitt-Str. 1, W-8014 Neubiberg/München. All brands or products are trademarks of their respective holders. ©1992. FORCE Computers, Inc.

SPARC. Don't let it slip away.

This is one of those times when you just have to seize the day. Because opportunities like this are rare indeed.

You see,




FORCE is the only company licensed by Sun to put SPARC[™] station 2 technology on VME.

Choose from a range of systems, for a highly integrated solution.

So we're the only ones who can give you validated hardware and software compatibility. Allowing you to run SunOS[™] with any of your SPARCstation 2 applications and peripherals. Without a hitch.

We also have a whole new family



Our CPU-2CE features a 40MHz SPARC RISC microprocessor, delivering 28.5 MIPS.

of SPARC 2 products. With everything from our CPU-2CE (6U) board to the *micro-force* single-slot system.

We even have 20-slot systems that allow you to customize I/O. Giving you the perfect combination of real-time and UNIX.

And our SPARC products provide the broadest software offering of any RISC architecture. Including SunOS, real-time operating systems and a host of application software.

So if you want a better grasp of embedded SPARC, call for a free brochure. 800-237-8863, ext. 5. Or in Europe, 49.89.608-14-0.

Because there's no reason to let all this power slip through your fingers.



CIRCLE NO. 6

OVI and VI get down to business

As special interest groups and committees study the feasibility of Verilog and VHDL working together, the bodies promoting the use of each HDL continue to operate as usual. Thirteen ASIC, EDA and Verilog end-user companies just elected Bill Fuchs, president and CEO of Simucad (Union City, CA) a Verilog HDL simulator company, to chair Open Verilog International (OVI). Meanwhile, vendor and user companies backing VHDL were making plans for the fall meeting of the VHDL International Users' Forum (VIUF), to be held later this month in Washington, DC. VIUF will convene under the leadership of VI president and CEO Michael Carroll and newly elected board chairman Harvey Jones of Synopsys. AT&T Bell Laboratories, CLSI Solutions, IBM, IKOS, Robert Bosch GMBH (Germany), Texas Instruments, and Toshiba (Japan) will be attending VIUF as new VI members.

In the meantime, a group of VHDL vendors and users is picking up efforts to generate support for the Cadence-backed initiative launched at DAC to incorporate standard Verilog modeling practices into VHDL, paving the way for the possible coexistence of the two HDLS. Known as VITAL (VHDL Initiative Toward ASIC Libraries), the initiative is now being coordinated by three industry representatives: Steve Schulz, member of the technical staff at the Defense Systems and Electronics Group at Texas Instruments (Dallas, TX); William Bilowitch, president of VHDL Technology Group (Allentown, PA); and Erik Hayskens, head of CAD and test for the VLSI group at Alcatel NV and the European coordinator for VITAL.

Being circulated among VITAL backers is a "discussion document" that's to be the foundation for a formal proposal to the appropriate VHDL standardization groups.

—Barbara Tuck

Intel deemphasizes Multibus

Over the past two years, Intel (Hillsboro, OR) has been relaxing its marketing efforts on behalf of

Multibus products and putting its energies into its PC-based lineup. Recent reports have the company's next move as removing Multibus products from the sales staff's commission schedule. Customers interested in buying Multibus hardware need not contact the Intel sales force; instead they should call (800) 438-4769 for any engineering, design, pricing, or other information. To place an order, you'll have to go to your local distributor. With no commission coming, it's unlikely a call to your Intel salesperson will result in much action.

With hindsight, it can be seen that Intel has been bowing out of the Multibus business since the departure of former marketing head Mike Richmond. Current marketing manager Dick Binns stresses that Intel is focusing its systems strategy on whatever will strengthen its preeminent position in the PC-compatible world. It can easily be conjectured from this that Multibus will play a secondary role to the PC business.

—Warren Andrews

Altera supports user-selected design tools

In a move to broaden support for its programmable logic families, Altera (San Jose, CA) put a strategic partnership program with tool vendors into place a few weeks ago. Dubbed ACCESS (Altera Commitment to Cooperative Engineering Solutions and Sales), the program promises the company's full support for whatever design-tool option partners and customers choose. Initial partners include Data I/O, Dazix, Exemplar, Logic Modeling, Mentor Graphics, MINC, and Synopsys.

"Altera recognizes the importance of supporting industry-standard platforms and design tools to enable our customers to create the design environments that meet the specific requirements of their projects," said vice-president of marketing Erik Cleage.

According to Michael Holley, senior staff engineer at Data I/O, "The significance of the ACCESS program is that it provides designers with the ability to utilize any of Altera's architectures in a known environment. In our case, ABEL

customers will now have support for all of Altera's broad range of current products, and in addition we will have the ability to introduce new device support coincident with production software from Altera."

The ACCESS program operates through framework integration and cooperative engineering involving the exchange of information in advance of product introduction. Design information may be transferred via EDIF netlist into MAX+PLUS II from design capture and synthesis tools, or exported to selected simulation environments.

—Barbara Tuck

Bus interface chips hit the street


Buscon saw the introduction of samples of some long-awaited bus-interface chip sets, and the announcement of some new ones on the way. Texas Instruments (Houston, TX) finally unveiled silicon of its Futurebus+ protocol controller, jointly developed with Force Computers (Campbell, CA). National Semiconductor (Santa Clara, CA) also announced availability of its Newbridge/National Futurebus+ protocol controller.

Force has announced a new partnership with Cypress (San Jose, CA) for the design and manufacture of a next-generation VME interface IC that supports VME64, SSBLT and many of the features in what was to be Revision D of the 1014 specification.

Not satisfied with the minimalist Futurebus+ ICs that were made available last month by TI and National, Futurebus+ board maker Cable and Computer Technology (CCT—Anaheim, CA) has huddled with LSI Logic and will be jointly developing a new Futurebus+ protocol controller. The CCT/LSI controller is expected to include many of the bells and whistles not included by the others. The chip will be the only one to fully support the 896.1 version of cache coherency and to include packet- as well as handshake-mode transfers. CCT has already built working boards with discrete (programmable logic) implementations of these functions, and antici-

Continued on page 12

Twice the Logic in Half the Space



DOUBLE-DENSITY™ FROM IDT

46% BOARD SAVINGS

IDT's new 16-, 18-, and 20-bit Double-Density FCT-T Logic family offers the performance of two octal logic devices in one flow-through 48- or 56-pin high-density, JEDEC-standard, shrink small outline package (SSOP) or Cerpack, for twice the functionality in half the board space.

A WIDEBUS™ UPGRADE

IDT's Double-Density logic family is more than twice as fast as ACT, uses 35% less power than ABT, and it's form-, fit-, and function-compatible with both

of TI's Widebus families. The Double-Density family also offers typical pin-to-pin skew of 250ps and quiescent supply current at 0.05mA (typ.).

3 APPLICATION CHOICES

5V High Output Drive

Ideal for low-impedance bus and backplane applications.

5V Balanced Drive (Low Noise)

Contains on-chip, source-terminating resistors to minimize signal noise. These devices are ideal for driving point-to-point transmission lines and highly capacitive loads, such as a bank of DRAMs or SRAMs.

3.3V Low-Power Logic

Designed for regulated or unregulated 3.3V power supplies, these devices use less power than 5V parts, without sacrificing high speed. 5V-to-3.3V unidirectional and bidirectional translators are also available.

FREE SAMPLES

Call today for free samples and a copy of the new **High-Performance Logic Data Book** and start your Double-Density logic design today!



| DOUBLE-DENSITY CONFIGURATION* | I _{OH} | I _{OL} | t _{PD} (Max.) | I _{CCQ} (Typ.) | PIN-TO-PIN SKEW (Typ.) | GND BOUNCE (Typ.) |
|-------------------------------|-----------------|-----------------|------------------------|-------------------------|------------------------|-------------------|
| High Drive | -32 mA | +64 mA | 4.1 ns | 0.05 mA | 250 ps | <1.0 V |
| Balanced Drive | -24 mA | +24 mA | 4.1 ns | 0.05 mA | 250 ps | <0.6 V |
| 3.3V | -8 mA | +24 mA | 4.8 ns | 0.05 mA | 250 ps | <0.3 V |

*Specs are for '244 device

Double-Density is a trademark of IDT. All others are trademarks of their respective manufacturer.

(800) 345-7015 • FAX: 408-492-8674
ASK FOR KIT CODE 3071



Integrated
Device Technology, Inc.

pates few problems in the chip versions. The CCT/LSI chips are expected to be available with both military and commercial screening grades.

In the meantime, there's a rumor that Motorola (Phoenix, AZ) may be working independently on its own Futurebus+ chips. The rumor has been fueled by the company's involvement with MOSAS and the greater interest it's expressed in Futurebus+. News is still sketchy about what kind of IC the company might be planning—and when. Recently, Motorola terminated talks with at least one system maker to jointly develop a chip. It's believed it did so because it plans to develop its own chip (a MOSAS module, perhaps?).

—Warren Andrews

Next finds big buy for 486 operating system

One of the little-noticed but potentially significant players in the current operating system wars is Apple cofounder Steve Jobs' Next Computer (Redwood City, CA). Next has been a marginal player in selling desktop computer systems ever since it introduced its first 68020-based machine almost five years ago. But inside the Next machine lies what may possibly be the big money maker—its NextStep operating system. Built around the UNIX-like Mach operating system from MIT, NextStep features a built-in graphical user interface and an object-oriented application builder.

Now Next has done the smart thing and ported NextStep to the Intel 80486 in the form of NextStep486, which is expected to ship early next year. That move has already attracted one very large potential customer, Chrysler (Highland Park, MI), which is reportedly ready to order 3,000 copies of NextStep486. Validation of the touted ease of development with NextStep in a large corporation such as Chrysler, which apparently wants to create its own mechanical CAD applications, could put another big player in the ring with those vying to become the heirs to UNIX.

—Tom Williams

X Windows on PCs to networks

One of the major benefits of the X Windows client-server architecture is that a user can run any application anywhere on a network from his or her own local terminal or workstation. Dedicated X terminals have traditionally come with X Windows and networking capabilities built in, but when a PC is to be the X platform, the user has had to install the X Windows software and usually the network card and its software and make the two work together.

Now an agreement between AGE Logic (San Diego, CA) and Novell (Monterey, CA) has resulted in PC-based X-Server software created by AGE, with integrated support for TCP/IP networking from Novell. The result is what AGE calls "push-button install" support for three environments—PCs running DOS, Microsoft Windows and Texas Instruments Graphic Architecture (TIGA)-based graphics controllers. The bundled X Windows/TCP/IP package will run on a PC with a variety of networking boards supporting the TCP/IP protocol. The package also supports the Open Datalink Interface (ODI) and the Internetwork Packet Exchange (IPX) for ODI, which will work with the AGE/Novell bundle support token ring and ARCnet networks as well as Ethernet.

—Tom Williams

Intel enters FPGA game

Intel (Folsom, CA) has announced its entry into high-density FPGA devices and development tools. Apparently Intel believes it can stake out a position in the FPGA market by offering simplicity—treating FPGAs like several PLDs on one chip. Intel's first device in its Flex-Logic FPGA family will be the iFX780. At the heart of the device are eight configurable function blocks. These blocks can be reconfigured either as logic or as memory. In a logic mode, a block has 24 inputs and 10 outputs, like a 24V10 device. In the memory mode, the block is configurable either as a high-speed SRAM or as ROM. These blocks are tied together by a global interconnect ma-

complex PLD routing limitations inherent in parts such as Altera's MAX and AMD's Mach," says David Stasaitis, product marketing manager at Intel.

The iFX780 also features some power-saving features. The non-volatile, programmable elements of the device are shadowed by SRAM elements. On power up, contents of the device's nonvolatile elements are loaded into SRAM. The nonvolatile elements are then turned off, resulting in lower power consumption. Running out of SRAM, only 1.5 mA per MHz of current is drawn. To put that in perspective, when the device operates at 50 MHz it draws only 75 mA. This makes the iFX780 a fairly complex device in a 132-pin package drawing less power than a single CMOS 22V10 (a 24-pin device). The shadowing SRAM technology also permits the iFX780 to be configured in-system while it's running. Designers can download new code during debug, and test the part without removing it from the board.

—Jeff Child


FPGA makers minimize handling

FPGA users are starting to think about in-circuit programming for devices mounted in fine-pitch PQFP. Data I/O's (Redmond, WA) director of software development, Dave Kohlmeier, says, "Manufacturers want to minimize handling when lead pitch gets down to 0.5 mm."

One FPGA maker who agrees is Lattice Semiconductor (Hillsboro, OR), which has applied its E2 process to a 2000-8000 PLD gate FPGA line called ispLSI.

Earlier this year, Cypress announced a flash 22V10, and Intel announced that a similar part was in the works. Presumably, either company's process could be extended to more complex devices.

—Don Tuite



THE ONLY ATTACHED PROCESSOR WITH FOUR ON THE BOARD AND 2.5 GIGAFLOP PERFORMANCE.

Mercury's MC860VS. The only attached processor that offers up to 32 Intel i860s in no more than 8 VME slots. So you get 80 Mflops to 2.5 Gflops of horsepower to handle demanding applications in defense signal processing and medical imaging. And all Mercury products can be configured in workstations and chassis systems to deliver scalable performance at a scalable price.

So if you're building or buying a high performance computing solution, take full advantage of Mercury's unparalleled investment in standards, innovative hardware, and software development environment. And get the most complete, flexible, high performance computing solutions available. If it's time you moved into the fast lane, it's time to call or write Mercury today for more information.

- ☐ Please send me more information on the MC860VS right away.
☐ I want to learn more. Send me your white paper "Multiprocessing for the 1990s" and enroll me in your Education Series program today.

Name _____
 Title _____
 Company _____
 Address _____
 City/State/Zip _____
 Phone _____

Please send to
 600 Suffolk Street,
 Dept. K,
 Lowell, MA 01854
 Fax (508) 458-9580
 (508) 458-3100

Computer Systems Inc.
MERCURY

CD 10/92

The Ultimate Performance Machine.

See us at DSPx Exposition & Symposium, Oct. 14-16, Booth #427

An Object Less For Absolutely Pr



The MACH™ Family From AMD: The Fastest, Most Predictable High Density PLDs Available Today.

Oops! You're a couple of nanoseconds shy this time, and it's going to hurt. Perhaps next time you'll choose a more predictable vehicle. And the most predictable high speed, high-density PLDs available are the MACH family from AMD.

Only the MACH Family offers you worst case delays of 15ns* or

less. Because MACH parts are essentially PAL® devices, just like the kind you already know. Not some hybrid PLD/FPGA, where you don't know how it performs — until it's too late. So you don't have to guess your delays or clock speeds, you just read them right off our datasheet.

But they're not just ordinary

| Model Number | Equiv. Gates | Macro Cells | Max. Delay | System Speed | I/O Pins | Hard-Wired Option |
|--------------|--------------|-------------|------------|--------------|----------|-------------------|
| MACH 110 | 900 | 32 | 12ns | 66.7 MHz | 44 | MASC 110 |
| MACH 210 | 1800 | 64 | 12ns | 66.7 MHz | 44 | MASC 210 |
| MACH 120 | 1200 | 48 | 15ns | 50 MHz | 68 | MASC 120 |
| MACH 220 | 2400 | 96 | 15ns | 50 MHz | 68 | MASC 220 |
| MACH 130 | 1800 | 64 | 15ns | 50 MHz | 84 | MASC 130 |
| MACH 230 | 3600 | 128 | 15ns | 50 MHz | 84 | MASC 230 |

*In applications with a full 16 product terms. Every MACH part is specified using real-life conditions with all outputs switching.

on In The Need edictable Speed.



PAL devices. They're bigger and better, with densities ranging from 900 to 3600 gates, all in our submicron CMOS technology.

Nor will you face unpredictable delays when you order. Because the entire MACH family is now shipping in volume.

Working with them is equally predictable. You don't have to learn any new techniques, just use the software and test equipment you already know. Like ABEL, CUPL, OrCad, and others. Not to mention the software and support from over 20 FusionPLD vendors — all prepared to bring your products to market on time.

And each MACH part can migrate easily to a pin-

compatible, hard-wired MASC™ counterpart for high volume. So you can get the volume you need, without redesign, NRE, or unforeseen delays.

So call AMD today at **1-800-222-9323**. And let the MACH family make your design cycle a whole heck of a lot safer.



Advanced Micro Devices

© 1992 Advanced Micro Devices, Inc. 901 Thompson Place, P.O. Box 3453, Sunnyvale, CA 94088. PAL is a registered trademark, and MACH and MASC are trademarks of Advanced Micro Devices, Inc. All brand or product names mentioned are trademarks or registered trademarks of their respective holders.

Can Your Current MCUs Meet These Specs?






comparable, high-speed MAC[®] counterpart for high volume. So you can get the volume you need, without redesign, NRE, or unexpected delays.
So call VMD today at 1-800-222-9822. And let the MAC[®] family take your design cycle a whole heck of a lot faster.



Advanced Micro Devices

For more information, please see the advertisement for the MAC[®] family in this issue of ESD: Design. Or contact your nearest Advanced Micro Devices sales office. For a complete list of sales offices, please see the back of this issue.

RAM devices. They're larger and better (ranging from 900 to 3600 gates, all in our standard CMOS technology).
Not will you face unpredictable delays when you order. Because the entire MAC[®] family is now shipping in volume.
Working with them is equally predictable. You don't have to learn any new techniques. Just use the software and test equipment you already know. Like ADEL, GUP, OrCad, and others. Not to mention the software and support from over 30 RISC/PC vendors — all prepared to bring your products to market on time. And each MAC[®] part can migrate easily to a 16-bit



Highest Integration

Lowest Power Dissipation

Low Voltage

Largest Memory

Smallest Packages



Your designs demand more from existing MCUs. As the leader in MCU technology, Mitsubishi offers you a breadth of 8- and 16-bit microcontrollers with the functions you want.

Highest Integration

Get the highest levels of on-chip functionality with a host of features you need most, including synchronous S I/O, UARTs, programmable I/Os, A-D/D-A converters, LCD and VFD drivers, PWM, special timers and more.

Lowest Power Dissipation

We offer the lowest power dissipation available today—4mA at 3.0V at 8MHz, for our most highly integrated 16-bit MCUs.

Low Voltage

For low voltage, our MCUs operate in the 2.7 to 5.5 volt range at 8MHz. For higher performance up to 25MHz, we offer 4.5 to 5.5 volt operation without paying the high power penalty.

Largest On-Chip Memory

Utilize up to 2 Kbytes RAM and 32 Kbytes ROM or OTP with our 16-bit MCUs, and up to 1 Kbyte RAM and 32 Kbytes ROM in our 8-bit MCUs, with expanded memory planned for the future. EPROM and ROMless versions are also available.

Smallest Packages


Get all these advantages in the smallest packages on the market. Mitsubishi offers a wide range of packaging options including 64- to 160-pin surface mount packages, with 0.50mm, 0.65mm and 0.80mm lead pitch. Shrink DIPs are available in 32- to 64-pin configurations.

For faster time to market, Mitsubishi provides development support tools for each phase of your product design.

For your best system solution, call Mitsubishi today at (408) 522-7498. For Eastern U.S. and Canada, call (408) 522-7499.



 **mitsubishi**
ELECTRONIC DEVICE GROUP



Your VME solution for real-time performance processing...

The Heurikon HKMIPS/V3500:

- 40 MHz R3500 Processor
- 8, 16, or 32 Mbytes EDC DRAM
- 64K each Data and Instruction Cache
- Intel 82596CA Ethernet Interface
- NCR 53C710 Fast SCSI Interface
- Up to 1 Mbyte Boot PROM


Real-time Solutions

HEURIKON
A COMPUTER PRODUCTS COMPANY

8310 Excelsior Drive • Madison, Wisconsin 53717
Phone 608 831 0900 • FAX 608 831 4249
800 356 9602

CIRCLE NO. 11

Looking for the best in a real-time multitasking kernel?



Since 1978, KADAK has provided products, documentation and technical support to those who absolutely insist on the best. It's one reason why more than 1,000 developers have come to count on KADAK for our AMX™ real-time multitasking kernel and more.

You'll marvel at the comprehensive, simple-to-follow AMX manual that will assist you in the development of virtually any application. You'll also appreciate the

prompt, courteous service you'll receive from KADAK product support people if you should ever need assistance.

All this, plus you will be given source code and automatic software updates for a full year after purchase. And, with AMX, there are no royalties to pay.

For further information on the KADAK AMX Real-time Multitasking Kernel contact us at our offices. Phone: (604) 734-2796 Fax: (604) 734-8114

Count on KADAK.



KADAK Products Ltd. Setting real-time standards since 1978.
206-1847 West Broadway, Vancouver, BC, Canada, V6J 1Y5

AMX is a trademark of KADAK Products Ltd. All trademarked names are the property of their respective owners.

CIRCLE NO. 12

COMPUTER DESIGN

Technology and Design Directions

EDITOR-IN-CHIEF/ASSOCIATE PUBLISHER
John C. Miklosz, Ph.D., (508) 392-2114

SENIOR EDITORS

Warren Andrews, *Computers & Subsystems*
(508) 283-2102

Barbara Tuck, *ASICs & ASIC Design Tools*
(516) 473-1661

Michael G. Donlin, *CAE/CAD Tools*
(508) 392-2123

Tom Williams, *Software & Development Tools*
(408) 335-5514

Don Tuite, *Integrated Circuits*
(415) 365-5656

Jeffrey Child, *New Product Developments*
(508) 392-2126

CONTRIBUTING EDITOR

Stephan Ohr, *Analog & Mixed-Signal*
(908) 232-1380

MANAGING EDITOR:

Arlyn S. Powell, Jr.

CHIEF COPY EDITOR:

Annette M. Staron-Wilson

COPY EDITOR:

Kym Wilson Gilhooly

PRESENTATION MANAGER:

Richard Sarno

EDITORIAL ASSISTANT:

Claire Ellis

ART DIRECTOR:

Jan Horner

AD TRAFFIC MANAGER:

Kelly Rice, (508) 392-2198

CIRCULATION DIRECTOR:

Robert P. Dromgoole, (918) 832-9213

CIRCULATION MANAGER:

Paul Westervelt, (918) 832-9287

REPRINTS:

June Bozarth, (918) 835-3161

PUBLISHER

David L. Allen, (508) 392-2111

ASSOCIATE PUBLISHER/ NATIONAL SALES MANAGER:

Tim L. Tobeck, (508) 392-2116

PennWell
PUBLISHING COMPANY

Advanced Technology Group

Dr. Morris R. Levitt, Senior Vice-President
Leslie P. Cypret, Vice-President, Administration
One Technology Park Drive
P.O. Box 990, Westford, MA 01886
Tel: (508) 692-0700 □ Fax: (508) 692-0525


Postmaster: Send change of address form 3579
to COMPUTER DESIGN Circulation Department,
Box 3466, Tulsa, OK 74101

ATTENTION RETAILERS: To carry *Computer Design*
in your store, contact International Periodical Distributors
at 1-800-999-1170, ext. 447 (dealer calls only).





Pick up the number one real-time operating system.



And run with it.

In today's competitive market, it's important to run with the best. And when it comes to real time, the iRMX® operating system is the clear favorite.

You see, iRMX has a 13-year track record for proven reliability. In fact, only DOS runs on more X86 systems. Now, iRMX for Windows provides the first real-time operating system with guaranteed response time that runs DOS and standard-mode Windows on the i386™ and i486™ architecture. That's right, true real time on a PC†

iRMX for Windows brings more than just

affordable hardware to real time. It also gives you a head start on development with access to the huge installed base of DOS applications and tools. In fact, you even have the option to run Windows†† (including Windows 3.1).


So take the first step, call (800) GET-iRMX (800-438-4769)* and ask for Lit. Pack. #2B. And start running real time with your favorite DOS and Windows software.

intel®

*Outside the U.S.A and Canada, FAX (503) 696-4633. © 1992 Intel Corporation. iRMX is a registered trademark and i386 and i486 are trademarks of Intel Corporation.

†iRMX runs on IBM and Compaq compatible i386 and i486 based PCs. ††All other products shown or mentioned are trademarks or registered trademarks of their respective owners.

Here's my thanks to everyone who helped put the Analog & Mixed-Signal program together.



John C. Miklosz
Associate Publisher/
Editor-in-Chief

We're missing the Halloween party again

A tradition started a few years ago at the PennWell Advanced Technology Group. When Halloween rolls around, we have an office party where all of the different publications and service departments go all out in a costume and area-decorating competition. As you might guess, *Computer Design* doesn't take competition lightly, and we've gone to great lengths to come out number one. Unfortunately, we couldn't take part in the competition last year, and we'll miss the fun again this year, because the **Analog & Mixed-Signal Design Conference** falls in the last week of October and too many of us are away when the costume and decorating work needs to be done. Missing the Halloween festivities and the laurels that we could gather drives us even harder to make sure that the Conference is truly world-class.

Well, October is here again and we think we've put together a great blend of tutorials, lectures and panel discussions for this year's Conference. You can find the details starting on page 57, so I won't reiterate here. What I would like to do is extend my personal thank you to all of the people involved in putting the program together. First, there's Ted Bahr of Miller Freeman, who first approached us with the offer to organize the technical program. It sounds like a cliché but Ted is a great guy to work with. Anastasia Kellow, John Huber and Chris McKim have also been great, and we learned a lot from them.

But the Conference wouldn't exist if it weren't for the many individuals and organizations giving their time to put presentations together. I don't have the room here to name everyone, but here's a shot at their company affiliations: AnaCAD, Analog Devices, Analogy, AT&T Bell Laboratories, Brooktree, Burr-Brown, Cadence Design Systems, Contec Microelectronics, Datel, Dazix/Intergraph, Deutsch Research, Elantec, Fujitsu Microelectronics, GEC Plessey Semiconductors, Gould AMI Semiconductors, Harris Semiconductor, International Micro Electronic Products (IMP), Linear Technology, LTX, Mentor Graphics, Micro Linear, MicroSim, National Semiconductor, NCR Microelectronics, Philips/Sigmetics, Quantic Laboratories, Racal-Redac, SGS-Thomson Microelectronics, Silicon Systems, Star Semiconductor, Texas Instruments, Viewlogic Systems, VLSI Research, and Zeeland Technology.

And who convinced these people that they should give us their time? The credit there, and our thanks, goes to Steve Ohr. You've all become familiar with Steve through his column in *Computer Design* every month, and the select few of you who'll be able to break away from your jobs for three days to attend the Conference will see another side of him when you get the chance to say hello.

And finally, I'd like to thank Patti Kenney, our technical programs coordinator, who took on responsibility for the myriad details that we consider "merely" details and who always reminds us of our responsibilities. We'd have a tough time pulling it off without her.

Synergy. Out Front Again... And Again.

The New Dual '040 VME SBC: Faster, More Efficient and Lower Cost Than Any Multi-board Solution

Get on-board multiprocessing and an astounding **40 MIPS** throughput, when you power your system with Synergy's new SV420 single-slot SBC. The SV420's dual-CPU design means you'll need fewer boards in your VME chassis, with no VMEbus bandwidth bottlenecks between the '040s, and system-wide cost saving of more than 40%.



And even if you don't need multiprocessing right now, the SV420 still puts you out front. Use the second '040 as a super-smart DMA controller. When combined with the SV420's **66 MByte/sec VME64®** circuitry, nothing communicates faster over the VMEbus.

Add even more on-board power by selecting from Synergy's big list of

high-performance, intelligent plug-on I/O modules – such as our latest Super-VSB module offering 40+ MB/sec over the VSB bus. Or choose another module, from a T1/E1 controller to a super-fast graphics engine, that plugs onto any Synergy SBC.

Better yet, just tell us what you need. We're the company you can talk to about your VME system design problems. You'll find that we listen *and deliver* (every Synergy I/O module on our list started as a customer request). We mean business when we say customer support is our most important mission.

So if you want to be out front in system performance, just be up front with Synergy. Call us today.




Over a dozen smart, powerful I/O modules fit Synergy's '020, '030, '040 and Dual '040 SBC's.



SYNERGY
microsystems

Synergy Microsystems, Inc.
179 Calle Magdalena, Encinitas, CA 92024
(619) 753-2191 FAX: 619-753-0903

We've just tested the performance of our new i960 CF processor.




Make no mistake. This is no middle-of-the-road processor.

It's our new superscalar i960™ CF chip. And it delivers blistering speed to your high-end printing, imaging, communications and networking applications. Simply put, it outruns every 32-bit embedded processor on

the market—and we have over 25 benchmarks to prove it.

You see, only our i960 RISC architecture brings superscalar processing with multiple operations per clock to embedded applications. Our new i960 CF processor is highly integrated with optimized data and instruction caches for throughput up to



twice that of our i960 CA chip. Which makes it the most advanced technology available for embedded applications.

If you're already using our i960 CA processor, you can quickly shift up to higher performance with full code and pin compatibility. And, of course, we offer a comprehensive array of Solutions960 development tools to


further speed up your design process.

So step on it. Call 800-548-4725 and ask for literature packet A9A50 for complete benchmarks. And learn what your applications can do with a faster engine.

intel®

CIRCLE NO. 14

ARTFUL *VME* SOLUTIONS



Advanced hardware is just the beginning.
Cyclone Microsystems completes the
development picture with C and Ada compilers,
VxWorks and pSOS+ real time environments,
source level debuggers, and a fantastic team
of support engineers, waiting to help
you create a masterpiece.

CYCLONE
MICROSYSTEMS

CVME961

CVME960

ISDN

Fiber

VSB

SCSI2

Serial

T1/E1

Ethernet

Greg Pope on: Software testophobia

While the United States is the unequivocal leader in the global software market, it's no secret that our software is infested with bugs. The existence of this anomaly is possible only in the absence of serious international competition. Like the automobile industry in the 1950s and 1960s, and the semiconductor industry of the 1970s and 1980s, the American software industry has used its pioneering innovations to sustain a prolonged period of unchallenged and prolific growth.


This period is now over, however. Foreign software developers are beginning to produce high-quality code, so the U.S. software industry must re-evaluate its entire development methodology to retain its competitive edge. If quality doesn't become a priority, the industry will suffer the same fate as the automobile and semiconductor industries.

The issue of software quality transcends the domain of American economic competitiveness and impacts everyone living in an industrialized society. As heirs of the technological revolution, we unwittingly interface with several million lines of software code everyday. Making a phone call, for instance, depends on the existence of over 500,000 lines of reliable software code. Digital alarm clocks, stereo systems, traffic signals, aviation equipment, medical devices, televisions, and automatic braking systems all rely on quality software programs.

The need for reliable, bug-free software, then, is no longer an issue that belongs exclusively in the software testing laboratory. It's an issue that affects our economic position in the global marketplace, as well as our everyday safety and well-being.

■ Why do we test?

The development of software applications and programs is a complex, arduous and tedious task, sometimes requiring several hundred thousand lines of software code for a single application. Because software development is still largely a human endeavor,



there will be errors. If these errors aren't detected and corrected, they can cause the software to malfunction. Depending on the specific application, a software malfunction can be a minor nuisance or, in the case of systems controlling hazardous materials, result in injury or death.

The demand for efficient and thorough software testing is magnified by the advent of CASE tools, software re-engineering, object-oriented programming, and other automated software development methods. These new tools and methods have let programmers generate an unprecedented amount of code with increasing rapidity. The result is a backlog of software that can't be adequately tested with today's debugging tools and manual methods. Now more than ever, there are profound needs for better automated testing methodologies and increased attention to the importance of software quality.

There exists in the American software industry a mythical six-stage software development process, or waterfall model, that's supposed to deliver an adequate project schedule. This schedule is predicated on several beliefs:

- System design requirements will remain frozen over the entire development period,
- Development resources will remain stable, and
- Software will be tested for up to half the period using state-of-the-art debugging tools.

At the completion of this development process, it's believed that the software will be bug-free and ready for market.

Greg Pope is the general manager of the test products group at Tiburon Systems, San Jose, CA.

In the real world, however, software testing is generally regarded as the stepchild of the software development process, rarely receiving the attention it deserves.

Why don't we test?

Based on the assumption that software can never be perfect, software management has come to accept the notion that incomplete or flawed software is accept-

able. This dangerous line of reasoning is often used by management to compress testing cycles and ensure a timely delivery to market. The result is often defective software and a continual devaluation of the software testing process.

Too often, software testing is relegated to second-class status in the development process because the system design is always more difficult and complex than originally planned. Additional code and significant design modifications are frequently required. Because the ship date must remain constant, delays in development frequently diminish the software-testing phase. The real-world software development process reinforces the all too commonly held assumption that software testing is inconsequential, an impediment to the development and financial goals of the software vendor.

The demands of a crowded and competitive marketplace require software vendors to deliver their products in a timely fashion. Because delays in ship dates can be devastating to a company's market share, management has come to accept the short-term solution of shipping incomplete and buggy software.

Beta testing and early adopter programs are often used as a substitute for thorough system testing, creating the potential for defects to go undetected. While beta testing was originally conceived as a method for evaluating the customer's software preferences, it has now grown into a bug detection methodology. The problem, however, is that a corporate

beta tester rarely has the time to significantly evaluate the product; when bugs are detected, the user usually gets frustrated and discontinues evaluating the product without forwarding the relevant information to the test engineer.

Defective software is knowingly shipped by most software vendors with the assumption that fixes and repairs will occur in the next version of the product. If software vendors are unable to make the necessary repairs in the current version, one wonders how they will find time to make them in the next version.

Yet the most insidious explanation for the prevalence of defective software is a complacent and undemanding software consumer. For whatever reason, users perpetuate the existence of software bugs by continually purchasing defective and incomplete products. We commonly accept the excuse that the "computer went down" as the scapegoat for what is really human incompetence or greed.

The short-sighted practice of shipping defective software isn't without its consequences. Increased demand for customer support, angry customers, costly post-ship bug fixes, loss of market share, decreased stock value, merger/acquisition, and bankruptcy are all

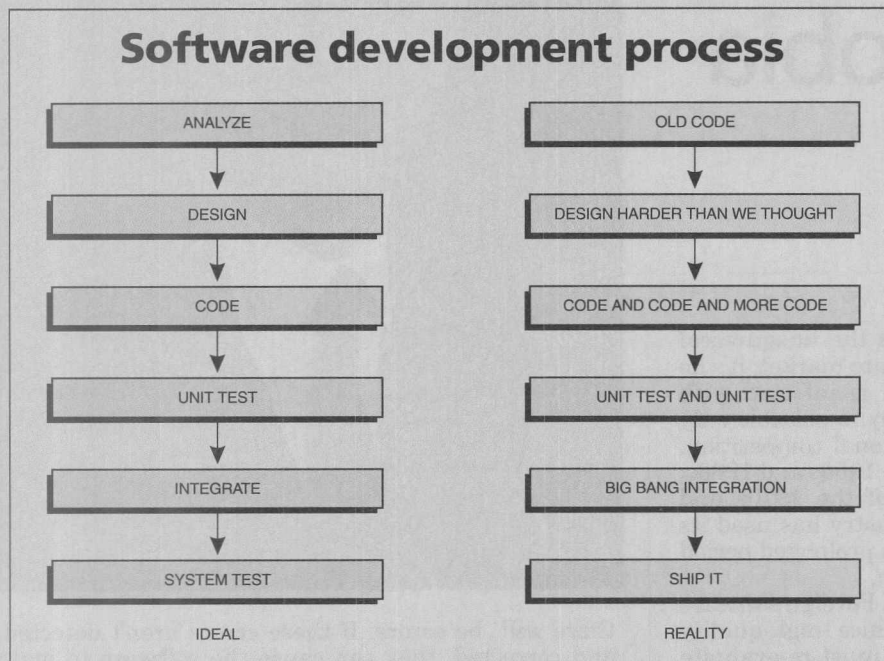
potential effects of adopting a software development philosophy that views testing as a dispensable impediment to corporate success.

Software "testophobia"

The software industry suffers from what I call "testophobia," an exaggerated, usually illogical fear of testing. Its symptoms include technical and business maladies. Test engineers who suffer from testophobia generally believe that bug repairs generate additional defects, creating a vicious cycle that perpetuates itself ad infinitum. Based on the myth that software testing can never be sufficient, these engineers experience a feeling of hopelessness.

Corporations whose software development teams suffer from testophobia experience late product deliveries, budget overruns, lost market share and profits, and a loss of user confidence. These symptoms, while not in themselves catastrophic, can collectively bring about the demise of any software vendor—witness the fall of Ashton-Tate.

Many of the causes of testophobia are rooted in an incomplete understanding of software testing methodologies and requirements. This misunderstanding proliferates among management because test engineers are perceived as a police force, a group that stands outside the software development process. The test group doesn't receive necessary management support because it's viewed as a bunch of spoilers that causes shipping delays.



comprehensive test designs and test tools are lacking. This factor, coupled with unrealistic test schedules and constant time-to-market pressure, creates a pessimistic test environment that often results in inadequate testing procedures and methodologies.

■ Cures for testophobia

Like other phobias, the only way to overcome testophobia is to desensitize the patient by slowly exposing him or her to the thing feared. With software testophobia, this means the entire industry must begin to actively confront the issue of software quality and its related testing assumptions, methodologies and tools. This desensitization process begins when the testophobic community acknowledges that the quality of American software is inadequate, and the only means to maintain existing global market share is to make testing an integrated part of the development process.

The software consumer must also become a part of this process. Consumers, on the whole, don't accept inferior, untested and incomplete products. It's unlikely that someone would purchase an automobile whose clutch was inoperable or a vehicle that won't include windows and upholstery until next year's model is available. Likewise, it would be inappropriate for a pharmacist to dispense drugs that haven't been completely tested. How about discounted airfares on experimental aircraft? When the consumer attitude toward software quality begins to change, the software vendors will make quality a priority.

Borland International (Scotts Valley, CA), a company that's made a concerted effort to produce high-quality software, has implemented a bounty game in which employees are paid significant cash awards for uncovering bugs in soon-to-be-released products. While this isn't an alternative to testing, it creates an environment that reinforces the importance of developing high-quality products.

Another solution that is currently gaining in popularity is harnessing the power of automated software-testing tools. If we use these tools to execute the more laborious and time-consuming tests, then test engineers will have more time to create the intelligent test cases that are necessary to produce better software.

The development of a legally binding definition of "reasonable" testing by industry experts and lawyers could provide consumers with a warranty that ensures the delivery of reliable, bug-free software. A related solution would be the creation of a center for software excellence, where all types of software can be independently tested and evaluated.

If we fear things in proportion to our ignorance of them, then it must be said that we are excessively ignorant about software testing issues. Fortunately, testophobia is not a permanent disorder. With a concerted educational effort, we can overcome our fear of testing and sustain our global dominance in the software market. Perhaps it's time to challenge ourselves to produce defect-free software within five years, before the offshore competition emerges with products that capture the lion's share of the market. There's nothing to be feared other than our own complacency.

NETWORKS.

Do they know something that you don't?



DYNAMIND™ 3.0

\$145

Everything you need to understand, build and use neural networks.

- Multiple training algorithms—Back Prop, Madaline III and TrueTime™ recurrent networks for forecasting
- Up to 8,000 inputs & outputs
- Up to 5,000 neurons per layer
- Extensive control of network parameters
- Remove Learned Patterns training option speeds up training
- Diagnostic, menu-driven interface
- Example-based User's Guide
- Mouse/keyboard access to On-Line Help

DYNAMIND DEVELOPER™

\$495

The complete solution for developing neural-network-based applications.

- Bundles DynaMind 3.0 with NeuroLink™, a library of neural network C-code routines
- Embed run-time neural networks
- Link multiple networks
- Create DFT neural networks
- Modularize complex problems
- Prototype the Intel 80170NX neural network chip
- No run-time license fee
- Compiles under Borland C 2.0 and Borland C++

Free Brochure

1-800-747-3531

NEURODYNAMX, INC. • P.O. BOX 323
BOULDER, CO 80306 • FAX 303/442-2854

* Sixth Generation Systems, January 1991 ©1992 NeuroDynamX, Inc. NeuroDynamX, DynaMind, DynaMind Developer, TrueTime and NeuroLink are trademarks of NeuroDynamX, Inc. Other products or brand names are trademarks or registered trademarks of their respective owners. All rights reserved.

CIRCLE NO. 16

LEISTER

Contact-free soldering and de-soldering

with the new Leister Hot-Jet "S" 

- For SMD and through board components
- Very rapid and ESD protected
- Hot-air temperature with stepless electronic control
- Air-volume with stepless electronic adjustment
- Power output 460 W, air volume 10-60 l/min
- Range of 700 desoldering nozzles available




ASK FOR FREE BROCHURE UW 218

Karl Leister, Elektro-Geratebau, CH-6056 Kagiswil/Switzerland
TEL. (041) 66 00 77, Fax (041) 66 78 16, Telex: 866 404

Brian R. White Co. Inc., 313 Henry Station Road
Ukiah, CA 95482 phone: (707) 462-9795
Farmingdale, NJ 07727 phone: (908) 938-2700

CIRCLE NO. 17

Synchronous 4Mb At 100MHz,



ASK FOR FREE BROCHURE UW 518

U.S. Microelectronics Corp., 315 Henry Station Road,
Farmingdale, NY 11737. Phone: (516) 255-2100.

U.S. Microelectronics Corp., 315 Henry Station Road,
Farmingdale, NY 11737. Phone: (516) 255-2100.

CIRCLE NO. 13

COMPUTER DESIGN, OCTOBER 1988

Cached DRAM. It Screams!

Matching low-cost DRAM technology with today's high-speed CPUs can be a design engineer's nightmare. Until now. **Introducing the 100MHz 4Mb Cached DRAM from Mitsubishi.**

FIRST SYNCHRONOUS DRAM

Mitsubishi combined a fast, 4K x 4 SRAM and a 1M x 4 DRAM with a wide, 16 x 4 bit internal bus and a synchronous clock design, all into one tiny TSOP IC. The result is the industry's first synchronous DRAM with on-board cache.

100MHz OPERATION

The Cached DRAM's large, 16 x 4 bit internal data path can transfer a 16-line data block in just one cycle, allowing the small on-chip cache to perform like a much larger external cache. The result is fast, 100MHz performance at a much lower cost than separate cache configurations. Plus, the Cached DRAM's fast copy-back scheme significantly reduces the miss cycle penalty time.

COST-EFFICIENT, SMALL SIZE

The Cached DRAM die and package are only 7% larger than those of a standard 1M x 4 DRAM. And, since they are manufactured with the same process and on the same production line as Mitsubishi's standard 4Mb DRAMs, Cached DRAMs are highly cost-efficient to manufacture.

LOW POWER OPERATION

With a clock that can be stopped to reduce power consumption to as low as 1mW, the Cached DRAM is ideal for portable and highly integrated applications where low power consumption, compact size and fast operation are essential.

MITSUBISHI'S CACHED DRAM PERFORMANCE

| Part Number | Cache Hit Access/Cycle | Cache Miss Access/Cycle | Direct Array Access/Cycle | Package |
|---------------|------------------------|-------------------------|---------------------------|---------|
| M5M44409TP-10 | 10ns/10ns | 70ns/280ns* | 70ns/140ns | TSOP** |
| M5M44409TP-15 | 15ns/15ns | 75ns/300ns* | 75ns/150ns | TSOP** |
| M5M44409TP-20 | 20ns/20ns | 80ns/320ns* | 80ns/160ns | TSOP** |

*Cache hit cycles can resume after one miss access time, while the copy-back completes in the background.

**TSOP Type II. Also available in reverse pin-out TSOP.


Samples Available

Not your ordinary next-generation DRAM, Mitsubishi's 4Mb synchronous Cached DRAM sets a totally new standard for cost-effective, high performance memory. For more information and technical specifications, please call (408) 730-5900, ext. 2106 or 2226.


Standard 4Mb DRAM.

Actual size.

4Mb Cached DRAM is only 7% larger than a standard 4Mb DRAM.



 **MITSUBISHI**
ELECTRONIC DEVICE GROUP



CCT is offering the first Futurebus+ -to-VMEbus bridge, comprising two boards: a 12-SU Futurebus+ card and a 6U VME card. The boards provide the essential link between the two buses, offering Futurebus+ systems much needed I/O. The bridge allows multi-cage distributed processing, shared memory structures between cages and Futurebus+ access to any existing VMEbus I/O devices. With the bridge, the Futurebus+ system can become a master on the VMEbus system, and vice versa.

First Futurebus+ -to-VME bridge board appears

Warren Andrews, Senior Editor

Futurebus+ is just starting to see some commercial success, but it's still far from critical mass in terms of available products—particularly I/O products. To fill the gap, Cable and Computer Technology (CCT—Anaheim, CA) has released the first commercial Futurebus+ -to-VMEbus bridge.

The company's timing is perfect. Digital Equipment Corporation (Maynard, MA) has already announced its Alpha-based workstation in a Profile B Futurebus+ format, which is expected to ship later this year. This may well be the first commercial system on the market, and purchasers will be hard-pressed to find any real I/O boards other than those supplied by Digital.

■ Some standard boards available

Thanks to the Navy's Next Generation Computer Resources (NGCR) project, there are a handful of standard Futurebus+ boards announced

or available, but they're CPU boards or other hardware designed expressly for military use. In addition, since the initial NGCR contract was awarded before completion of the Futurebus+ specification, it calls for the original soft metric form factor, leading to boards that won't work in a 12-SU rack.

This potential deficiency in Futurebus+ I/O boards, however, was anticipated early in the definition of the bus, and a bridge specification has been in the works for more than a year. This document is now in draft stage, as 1014.1. In addition, it's been Digital's strategy from the beginning to provide users with I/O through a bridge to another bus—specifically VME.

Digital has had little to say about making deals with third-party developers of bridge products like CCT, but it has certainly been in the forefront when it comes to sampling existing boards. Early hints of this interest could be seen when Digital tried most of the existing Future-

bus+ cards in its own backplanes. While the company hasn't published any results of the tests it conducted, it did indicate it was able to get CCT's 68030 processor board up and running in minutes. No statement was made about the other boards tested.

"The Futurebus+ -to-VMEbus bridge uses the same interface as the single-board computer we tested [CCT's 68030 board], and we have complete confidence in the implementation of the bridge product [made by CCT]," says Digital's Futurebus+ marketing manager, Steve Justus. This means Digital will be able to offer VME-based I/O to its customers. But perhaps more important, says CCT product line director of Futurebus Bruce Kimble, "It clearly shows the soundness of the Futurebus+ specification and demonstrates that manufacturers working independently of each other can produce boards that work together on the same backplane."

While Digital is gearing up to ship its Futurebus+ -based workstation, CCT is providing the first Futurebus+ -to-VME bridge board, promising a link between existing VME crates and the emerging world of Futurebus+. Even before the bridge specification has been completed, CCT has released its FBB-001 interconnection strategy.


The bridge specification, originally part of the Revision D upgrade of the 1014 VMEbus definition (on its way to becoming 1014.1), isn't expected to surmount the final obstacles in the standardization process until late this year or early next, but CCT's Kimble believes the existing document is tight enough to design a product around now. If minor changes appear in the specification, he says that they can be accommodated with only small changes to the board set.

■ A and B connections

CCT's bridge lets any Futurebus+ Profile A or B system connect directly to standard VME devices. Either the Futurebus+ or VME system can be configured as a coprocessor in the overall design, or either can serve as the only processor on the other system. No other processor or system controller is required to make CCT's product work; the VME board that's part of the bridge performs the Slot 0 controller functions.

The bridge comprises two boards,

NEW & NOW



*Ironics IV-3234
Very Intelligent SIO*

40 MIPS SIO, 16 Ports, Free '030


Finally, a smart SIO board that puts the intelligence where it really counts: on the serial communications controllers. *A full 40 MIPS worth!*

The result? *Processor power to spare!* The on-board 68030 is free because each channel has its own DMA and 16 byte FIFO controlled by a RISC core. With all 32 channels operating simultaneously at 64-Kbits/second, over 95% of the main CPU's bandwidth is still available. Now that's intelligent!

The Ironics IV-3234 is not only smart, it's flexible. Standard features include *synchronous and asynchronous modes*, 1 or 4 MBytes of fast DRAM, and the most powerful VME interface in the industry. SCSI/Ethernet, VSB, and additional higher speed SIO ports (up to T1 rates) are optional. Choose from among the *best* development environments and real-time kernels.

So make the smartest choice for intelligent SIO...the Ironics IV-3234.

...and now, *more SIO in less space with no hassles!* Ironics' new Modular SIO system fits 16 ports, with 8-signals and modem controls on every port—in *just one VMEbus slot!*



*Introducing Ironics
MODULAR SIO—the most elegant,
comprehensive Serial I/O protocol conversion
and signal distribution system in the world*

*Call our VME Product Hotline Today.
1-800-334-4812*



IRONICS INCORPORATED • 798 Cascadilla Street Ithaca, New York 14850 USA • 607-277-4060 Fax 607-272-5787

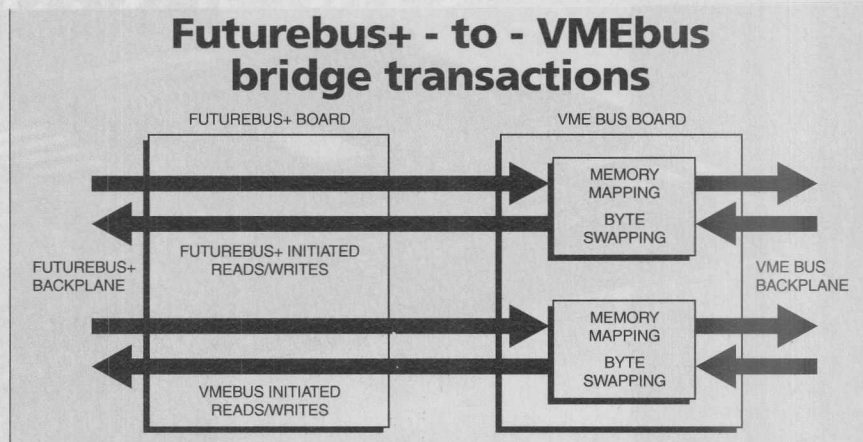
IV-3234 and Very Intelligent SIO Controller are trademarks of Ironics, Inc. All other brands or products are trademarks of their respective holders.

one for the Futurebus+ backplane and another for the VME system. The boards are interconnected with a round, 100-conductor, EMI-shielded cable which lets the two systems be separated by as much as six feet.

Windowing technique

In operation, a memory windowing technique lets you move information from one system memory to the other. The Futurebus+ system can access memory and I/O in the VME system, for example, while a VME master can access Futurebus+ memory and I/O.

The windowing technique works by mapping memory addresses from the target system's address space onto the initiator's address space. The window size and location is programmable through a set of page-map registers located on the bridge boards. The bridge permits flexible memory mapping, using a 4-kbyte page size with 8-kbyte page-map registers for Futurebus-to-VMEbus transfers and 32-kbyte page-map




CCT's Futurebus+-to-VMEbus bridge maps a memory address from one bus to the other. Inherent in the memory mapping scheme is the ability to swap bytes around when different processor types are used. The bridge also permits multiple crates to be tied together in a star or daisy-chain configuration.

registers for VMEbus-to-Futurebus+ transfers. The memory-mapped architecture also allows byte swapping when needed.

When a processor performs a ran-

dom-access read or write to an address within the window, the read or write is translated by the bridge board from its address space into a read or write on the target system.



RISC

Solutions

Your VME solution for real-time integer processing...

The Heurikon HK80/V960E:

- i960CA delivers superscalar performance for applications requiring fast data movement
- DMA-supported I/O such as Ethernet, SCSI & 4 serial ports
- VxWorks[®], pSOS[™] or C Executive[™] speeds your development
- Complete system integration services available
- Ask about Heurikon's other RISC & CISC CPUs

Real-time

HEURIKON
A COMPUTER PRODUCTS COMPANY

8310 Excelsior Drive • Madison, Wisconsin 53717
Phone 608 831 0900 • FAX 608 831 4249
800 356 9602

CIRCLE NO. 20




reprints

Reprints of any article or advertisement appearing in *Computer Design* may be ordered from: June Bozarth, CSR Reprints, **PennWell Publishing Company**, P.O. Box 1260, Tulsa, Oklahoma, 74101, 1-800-331-4463 or 918-835-3161 Ext. 379.

Minimum order: 100 black and white copies, 500 four-color copies.

**COMPUTER
DESIGN** *Technology
and Design
Directions*

CIRCLE NO. 21



TI's 7.5-ns '22V10. For those times you're torn between profit and performance.

Trying to get the best performance on a tight budget can create a few hang-ups. Texas Instruments has a simple solution. Introducing the TIBPAL22V10-7. TI's newest high-performance programmable logic device that's designed to fit the bottom line as easily as it meets your design specs.

High performance, low price

While competitive pricing is one of our '22V10's most outstanding features (less than \$17 when you purchase 5,000 or more), you'll be even more impressed by its performance.

At an incredibly fast 7.5 ns, our '22V10 supports system speeds up to 50 MHz with a variable term distribution that gives you more design freedom with complex functions. It's an excellent choice for high-end systems

using the latest microprocessors. And since all this is achieved using our proven bipolar process, the '22V10 provides a universal architecture that's easy to work with using familiar design tools.

Accurate, dependable and available today

Speed and ease of use mean nothing if difficult programming keeps your product from getting to market on time. That's why our '22V10 is designed for quick, dependable programming with your present tools. In fact, we're running at a 99.4% first-time programming success rate.

Best of all, our 7.5-ns '22V10 is available in volume *today* with just-in-time and ship-to-stock delivery programs tailored to meet your needs. You'll also have the backing of our

global support network to help keep things running smoothly.



Hang in there — a free sample is on the way
Simply return the attached reply card or call 214-995-6611, ext. 3717, for one free TIBPAL22V10-7.

 **TEXAS
INSTRUMENTS**

A built-in DMA processor lets the system perform block data moves from one bus to the other. Using event channels, interrupts can be passed between the two systems.

CCT's bridge lets you connect multiple Futurebus+ and VME systems in star or daisy-chain configura-

tions, or in some combination of the two topologies. Software included on each bridge card takes care of power-up initialization, self-test and default configuration. If you add an RS-232 terminal to the port provided on the Futurebus+ module, diagnostics can be performed to exer-

cise the functioning of the bridge. You can also perform minor I/O tests on the buses.

Independent operation

A key feature of the FBB-001 is that it's not just a repeater connection linking the timing of both buses, in which case activity on one bus would slow down the other. Instead, the FBB-001 permits each bus to operate asynchronously and independently of the other. The two buses are linked only when a memory or I/O reference is made to an address that translates the reference to the other system.

In addition, the Futurebus+ module fully supports split transactions, an operation that transfers data with the Futurebus+ system running at full speed while still accessing data on the slower VMEbus. The Futurebus+ card also supports central arbitration in the Futurebus+ system.

At a time when the trend is toward smaller and smaller I/O boards, such as SBUS, Turbochannel, PC-104, and even IndustryPacks, it seems unlikely that many independent Futurebus+ cards will be developed.

Further, with the advent of the Modular Open Systems Architecture Standards (MOSAS) Committee there's a growing commitment to module standards, and Futurebus+ CPU may always be dependent on bridges or mezzanine approaches for I/O. This, of course, raises the question of the value of the B (or I/O) Profile of Futurebus+. Will it be relegated to supplying a bridge function and carrier boards?

Such a scenario may not be that far off. CCT's Kimble has hinted that CCT is considering a Turbochannel-to-Futurebus+ bridge. And while Digital continues to play its cards close to the vest there have been many hints that the company plans to double the 80-Mbyte/s transfer rate of Turbochannel so it can keep up with Futurebus+-based workstations. ■

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

Cable and Computer Technology (CCT)
(714) 937-1341 Circle 201
Digital Equipment Corporation
(508) 493-6293 Circle 202



Why Settle For Second Best When You Can Have the Realtime, Real UNIX®?

VenturCom provides VENIX™ operating systems for realtime and embedded realtime solutions on popular low-cost Intel® 386 and 486 processors. VENIX software is the enhanced, USL licensed, UNIX operating system—not a look-alike or a clone.

VENIX System Software

- Full Kernel Preemption
- Deterministic
- Priority Scheduling
- ROMable/Embeddable
- Contiguous File System
- 27 µs Interrupt Response

Call for our Realtime Technical Overview.

VenturCom
Industrial Strength UNIX

215 First Street, Cambridge, MA 02142
TEL 617-661-1230 FAX 617-577-1607

A Truly New Compiler Comes Along Only Once In A Generation.

Introducing **Ultra C**™. It lets you explore the boundaries of performance for hard real time.



When performance counts, you need a C compiler that will produce the fastest, tightest code possible—especially in demanding hard real-time applications. Yet most C compilers commonly used today were originally written ten or more years ago. Since then, computer scientists have learned a lot about how to make better compilers.

That's why Microware undertook the epic task of creating an all-new C compiler based on the latest academic research—a compiler able to eke out every drop of performance from modern 16-, 32- and 64-bit CISC and RISC microprocessors. Microware's **Ultra C** uses a modular architecture that creates a virtual playing field for vigorous action by dozens of optimizers to analyze, arrange, accelerate and compress code into fast, compact executables. Heuristic analysis of register and variable usage *far surpasses the capabilities of what most human programmers could do themselves.*

State-of-the-art, plus...

Of course **Ultra C** has all features expected in a state-of-the-art compiler, such as full ANSI C compliance, C source code symbolic debugging, and a comprehensive set of standard libraries including support for the extensive real-time capabilities of the OS-9® and OS-9000® Real-Time Operating Systems. Options allow selection of either ANSI or K&R compatibility. You can even turn the optimization knobs yourself to best match your application's requirements.

Built-in quality and reliability

You can count on **Ultra C** right now, because it's probably the most thoroughly tested new compiler in history. Before the first copy went into beta test, it successfully passed the massive Plum Hall ANSI C Validation Suite—eight major sections, in all over 500,000 lines of C compiler torture track. Compilers are not a sideline at Microware—we've been developing them in-house since 1978.

CISC now, RISC soon

Ultra C is available now for all 680X0 family and 386/486 family CPUs running Microware's OS-9 and OS-9000 Real-Time Operating Systems. RISC versions are coming soon.

Call Microware Today!

1-800-475-9000

In California, call (408) 980-0201

microware

MICROWARE SYSTEMS CORPORATION

1900 N.W. 114th Street • Des Moines, Iowa 50325-7077


England / Benelux / Scandinavia: (44) 703 601990 • Germany: (49) 6221-862091

France / Spain / Italy: (33) 42.58.63.00 • Switzerland: (41) 56-83-3377

MORE CHOICES • MORE OPTIONS • TOTAL SUPPORT

CIRCLE NO. 24

Mass storage



without the mass.

That's AT&T "Customerizing."

A 3-chip disk drive solution!

A mass storage chip set that eliminates as many as 5 to 8 ICs as well as a dozen discrete components — drastically reducing your time-to-market while maximizing your disk capacity.

That's what "Customerizing" is all about!

Low power solution

AT&T's solution sharply decreases disk drive power demand, increases battery life, and allows for lower heat dissipation, a critical concern in small form-factor devices.

It also supports multiple data rates for zoned band recording as well as embedded servo formats and split data field techniques.

All-CMOS solution

The chip set is designed in a common CMOS process, allowing for future integration to a single device, resulting in even greater space and power savings, and producing the smallest possible form factor.

All the devices are highly

programmable, permitting product differentiation and evolution — and providing unprecedented levels of power management.

Integrated solution

Our hard drive solution features the first all-CMOS read channel ICs, fabricated with AT&T's advanced, low-power 0.9 μm drawn mixed-mode CMOS process.

This CMOS process advantage is combined with high integration and substantial system design and test expertise to minimize the number of external components needed to implement disk drive electronics.

The chip set is supported by evaluation kits, which provide both a hardware and a software platform on which to build a prototype disk drive, as well as a fast means of evaluating the devices.

To get the evaluation kit, along with complete technical information on AT&T's 3-chip disk drive solution, just give AT&T Microelectronics a call at 1 800 372-2447, ext. 908.



AT&T
Microelectronics

MOSAS wants military to drive commercial standards

Warren Andrews, Senior Editor

A consortium of major defense contractors is trying to shift all military electronics to a mandated standard platform. This platform would be based on existing or emerging standards developed by the IEEE, ANSI, JEDEC, and other standards bodies. An implicit part of the plan is that the same standards and components will be used in both military and commercial systems.

The group, known as MOSAS (Modular Open System Architecture Standard), has been gaining momentum over the past several months, but has kept a low profile. Knowledgeable non-members are reluctant to talk on the record about the group, fearing retribution from customers among the country's largest defense contractors, 20 of whom belong to MOSAS.

■ Suggested standards

While the group hasn't yet specified all the standards it will endorse, it seems to be committed to the Futurebus+ 896.1 logical layer, and to the SEM-E (Standard Electronic Module—E Format) as the standard board form factor. The SEM-E card is approximately 6 × 6 in., or slightly smaller than existing 6U VMEbus cards. Because of its relatively small size, SEM-E has been the unchallenged choice for avionics systems.

The MOSAS group has also settled on the 394-pin Teradyne connector over the 2-mm hard metric connectors defined by the standard Futurebus+ profiles. The 2-mm Metral-type connectors couldn't meet military muster, it seems, while the Teradyne unit is already the high-density connector of choice in a number of military systems.

MOSAS has also proclaimed the need for a modular approach to systems, calling for smaller multi-

chip modules (MCMS) to be included in the standard. As part of its MCM effort, the group has apparently singled out a JEDEC surface-mount package with 343 pins. The JEDEC standard also defines limits for the physical form factor, letting four MCMS reside on one side of a SEM-E card. Although it's been suggested that surface-mount technology would let eight modules be attached (with four on each side), difficulties in handling I/O and dissipating heat haven't been resolved, so this isn't currently possible.

Sources who've declined to speak publicly have speculated that some standard modules have already been defined, including a 68000-based processor with memory, a dual 96000-


based DSP module with fast static RAM and a Futurebus+ interface module. It's also been suggested that an 88000-based module has been adopted. From this it can be inferred that Motorola, sole maker of the 68000, 96000 and 88000, is one of the participating MOSAS members.

■ What's different

MOSAS contends that its goal is to standardize off-the-shelf modules and boards so that military applications can take advantage of the economies of scale offered in the commercial and industrial markets. "The military has been using very similar modules and systems for several years," comments one MOSAS supporter, "only each module and board was proprietary to only a single system. Because each module was custom, a single one could cost \$10,000, \$15,000 or more. An equivalent commercial or ruggedized industrial module, if available, would cost no more than a few hundred dollars."

A further goal is standardizing all military computers, potentially leading to far greater economies being realized in areas from acquisition to maintaining spare parts inventories, and from technical training to easy line and shop board and module serviceability.


In its recommendations to the Under Secretary of Defense (Acquisition) (USD(A)), the MOSAS council suggested that three directives be issued. The first is a re-issue of the DOD directive concerning the "development and use of nongovernment standards developed by nongovernment standards bodies." A second directive would establish a policy requiring the use of a modular, open-system architecture standard, developed by nongovernment standards bodies, that incorporates IEEE, ANSI, NGCR, and MOSAS open-system specifications that ensure design, development and upgrading of weapon systems, command and con-



The MOSAS council is suggesting standard board and module formats for both military and industrial applications. For individual modules, MOSAS has singled out a JEDEC surface-mount pinout with 343 pins. This MCM is called a shop replaceable module. The 6-×-6-in. SEM-E board is small enough that four MCMS can fit on one side; it's called a line replacement module (LRM). The line replaceable unit (LRU) shown here provides a backplane, chassis, card cage and interface electronics; LRMs fit in any card slot in the backplane, since interoperable, plug-compatible SEM-E cards are used.

Radstone reveals an ancient technique for designing an ultra high performance VME system...now.

Lurking right there in the VMEbus designer's bag of powerful tricks is the overlooked, under-utilized and very potent VSB with all of its board and system enhancing features. There's even a complete range of VME/VSB boards just waiting to be designed into a superfast, super-flexible system...right now.



A multiple systems-within-a-system architecture forming tightly coupled sub-systems capable of independent and parallel operation.

The VME Sub-system Bus gives you...

- A proven system solution that's multi-vendor sourced and industry accepted.
- Highly flexible multiple bus capability today.
- Multiple systems within a system.
- Wide range of products: CPUs, I/O and memory.
- Freeflow+ architecture that eliminates system bottlenecks.
- Increased system efficiency: VSB for data/VME for message passing.

If you're looking for a way to get *next generation* bus architecture performance right now, look at VSB. Radstone will support your efforts with the world's most complete line of VME/VSB processor, I/O and memory boards and the world's *only* line of Freeflow+ architecture boards. Call us to take advantage of the ancient wisdom of VSB...now.

Radstone Technology Corporation
20 Craig Road, Montvale, NJ 07645-1737
Call Toll-Free: (800) 368-2738
Eastern Region: (201) 391-2700
Central Region: (708) 304-0202
Western Region: (408) 727-4795
CIRCLE NO. 27

R
RADSTONE
TECHNOLOGY

COMPUTERS & SUBSYSTEMS

trol systems and information systems so as to provide true interoperability across systems by using common, plug-compatible electronic modules. A third directive from the USD(A) would promulgate and implement a three-service, open-system architecture approach, such as has previously been discussed among the service branches.

■ MOSAS isn't all roses

It's difficult to fault MOSAS' desires to improve embedded military computers while significantly cutting taxpayer cost. It's been suggested, however, that there may be a few flaws in its proposed implementation.

First, while the military establishment is a large user of computer and electronic hardware, it's dwarfed by the industrial and commercial markets, domestic and foreign, so it's unlikely that military standards will drive these larger

markets. Also, one laudable goal of MOSAS is to use COTS (commercial off-the-shelf) and NDI (non-developmental item) hardware, which will emerge from the commercial sector. If standards are to evolve, then, they will—and should—come from the commercial sector, and not MOSAS.

The connector selected by MOSAS, for example, is a militarized device, totally different from the 2-mm connector chosen for the Futurebus+ specification. At this time, it's also a single-sourced connector, and no one making the 2-mm connector offers a militarized version, the feeling being that the volume of business to be gained isn't sufficient to retool the commercial version.

Second, not every product type will have a common denominator. While it's a comfortable thought that every system can take advantage of the SEM-E form factor and its attendant modules, for example, the

bottom line is that such approaches will probably not be effective across the board. Certainly, because of size and weight restrictions applying to airborne electronics, the SEM-E approach is valid for avionics systems, but shipboard and land-based systems don't have the same requirements. A move to that one form factor would undoubtedly push up the cost of commercial products beyond normal market elasticity.

■ Cost considerations

Third, the use of standard modules and a standard local bus might have an adverse effect on cost compared with other standard approaches, depending on the local bus used. To develop a standard bus with all the bells and whistles needed to handle cache-coherent processor modules, for example, could raise the cost of module interfaces prohibitively. On the other hand, should the local bus be strictly a peripheral bus, it

We Make 41% Of The World's Memory Cards.

With a 41% market share,* Mitsubishi is the world's leading supplier of memory cards.

The largest share of today's memory card users depend on Mitsubishi for two good reasons: meticulous service and the best selection.

We manufacture cards in all five memory types — DRAM, SRAM, EEPROM, flash EEPROM, OTPROM and MROM — in the highest densities (up to 16Mbytes) and in the current versions of PCMCIA, JEIDA and

JEDEC standards. Plus, we offer custom cards and custom panel artwork. We can mix memory types, consolidate logic into ASIC, even add MCUs on board.

Whatever it takes, we'll work with our customers to achieve their memory card needs. With engineering and marketing support, well-stocked inventories and automated shipping services all onshore, Mitsubishi is the world's leading memory card source.

Call (408) 730-5900, ext. 2214.



MITSUBISHI
ELECTRONIC DEVICE GROUP

*Based on independent market research. Information available upon request.

Rugged-MIL VMEbus: Setting New Standards

■ Small in Size – Great in Performance

or's 3U boards offer more CPU power and functionality on less than half of a 6U board.

■ Low Power CMOS

The 100% CMOS design guarantees lowest power consumption: e.g. our 68030 CPU board requires only +5 V, 0.85 A at 32 MHz!

■ -55° to +105° C

The cold CMOS design combined with conduction cooling assures operation at the extreme of +105° C.

■ 100 G Shock

or's stiffened 3U boards are almost completely shock and vibration proof. Jumpers and sockets have been replaced.

■ Highest MTBF

or's ruggedized boards have set marks in VME: e.g. our 68030 CPU boards has an MTBF of 740,000 hours! (MIL-HDBK-217, 20° C).

■ Broad Product Range

68040 CPU, 68030 CPU, dual 68302 communication contr., graphics, MIL-STD-1553B, ARINC-429, SCSI, ethernet, flash-disk, synchro-resolver, digital I/O, analog I/O, ... VRTX-32, PSOS+, OS-9, ADA, C, ...

■ 4 Styles – Low Cost

Four software-compatible styles from economy to MIL-883/B cover all environmental needs at best prices.

■ Controlled Quality

or's quality assurance complies with AQAP-4 and MIL-I-45208.

or

Industrial Computers
Sieglingenstr. 19 1/2, D-8900 Augsburg 1
Tel +49 821 5034-0, Fax 5034-119

In the US contact: **DYNATEM**
15795 Rockfield Boulevard, CA 92718,
Irvine, Tel (714) 855-3235, Fax 770-3481

Rugged Laptop,
Waterproof,
NEMP protected

Conduction Cooled Housing
with optional Liquid Cooling



CIRCLE NO. 29

wouldn't work for multiple, cache-coherent processors.

If MOSAS selected a standard local bus such as PCI, recently introduced by Intel, then standardized, inexpensive, readily available modules for memory and I/O could be used. But the current MOSAS strategy


seems not to include selecting one of the existing or emerging local buses. An alternative would be to develop yet another standard exclusively for the military, since PCI and other local bus approaches, such as VESA, already have followings in the commercial market. With this alterna-

tive, each module would be developed specifically for an application, resulting in little or no savings from current approaches.

Fourth, there's little indication that a module standard created for a handful of military applications will find much acceptance in the commercial market. Economies of scale won't be realized in such a case. It's far more likely that commercial systems will be modified for military use. A modified version of a commercial Hewlett-Packard workstation, for example, has been adopted for shipboard use as a ruggedized tactical computer, providing 80-Specmark performance for \$28,000—less than the cost of a single card or module in a full MIL-STD system.

If the intent of MOSAS is to keep the military electronics business proprietary and restricted to an insider group of companies, as some critics maintain, then the effort will run into problems. But if MOSAS indeed aims at opening up military systems to industry-wide standards, then the group has done little to earn criticism. At this time MOSAS has no charter to develop standards from any official standards body, such as ANSI or the IEEE.

But, as one MOSAS supporter says, "The military contractors have been doing things the same way for a long time. It has to be recognized as a major change that they are considering any open standards at all—no less those developed by commercial developers." Add to this situation the fact that military systems development has always been done secretly, and it's little wonder that the major players are reluctant to lay their cards on the table at this early stage of their policy-making. ■



320 MFLOPS

SuperCard i860-based vector processors deliver. With our Quad-860 multiprocessor and our SC-3XL ultimate performance models, CSPI continues to offer you the greatest range of vector processing power. Look what you can get on a 6U board with VME64:

- 320 MFLOPS Performance
- 64 MBytes Memory
- 200 MBytes per second I/O
- Intel's 50 MHz i860 XP chips
- VME64 and VSB
- MIL Spec
- GFLOPS Multiprocessors
- 200 usec 1024 complex FFT

CSPI's SuperKit and Integration Toolkit software cut development time to the bone with:

- Industry-leading Libraries
- pSOS+ Real-time Kernel
- UNIX compatibility
- DataFlo Visual Programming
- C, Fortran, and Ada

Call or write CSPI and learn more about our new Quad-860 and SC-3XL SuperCard family members: CSPI, 40 Linnell Circle, Billerica, MA 01821, 1-800-325-3110, Fax: 508-663-0150. In Europe, CSPI, Creteil, France, 33-1-43774577, Fax: 33-1-43994508.



CSPI

CIRCLE NO. 30

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

A.C. Stensrud
Chairman of MOSAS Council
(602) 948-1634 **Circle 203**

Al Marcantonio
Vice-Chairman of MOSAS Council
(401) 847-8000 **Circle 204**

BECAUSE COMPUTERS WERE MEANT TO RUN TOGETHER.

Maybe it's not natural to expect a network of microcomputers to perform like a supercomputer.

With the QNX operating system, it's not natural to expect anything less.

POSIX AND MORE

Thousands of VARs and OEMs choose QNX for mission-critical applications – from POS to manufacturing to medical instrumentation.

And with good reasons.

Like POSIX compliance.

And *real* realtime performance.

And true microkernel architecture.

And message-passing IPC.

Not to mention our technical support and customer services.

Now add "FLEET" to the list.

FAULT-TOLERANT NETWORKING

With most networks, a hardware failure spells disaster.

But not with QNX.

If a card or cable fails on a dual-net FLEET setup, QNX will automatically re-route data through the other network before you – or your application – can even blink.



LOAD-BALANCING ON THE FLY

For greater throughput, the FLEET network puts all available network hardware to work *at the same time*.

And it will dynamically distribute the load by choosing the best route for the job.

EFFICIENT PERFORMANCE

FLEET uses network hardware to full advantage for maximum throughput. Whether you're running Ethernet for speed (application-level throughput at just under 1 Mbyte per second) or Arcnet for deterministic transactions – or both network cards in the same machine – you can count on FLEET for optimum efficiency.



REALTIME OPERATING SYSTEM

EXTENSIBLE ARCHITECTURE

You can support new networks simply by adding new drivers.

And you can start and stop drivers *dynamically*, without even rebooting.

TRANSPARENT DISTRIBUTED PROCESSING

In QNX there's no difference between local execution and network-remote execution.

Which means you don't need to modify your applications in order to distribute them across the FLEET network.

To sum up: networking with QNX is fault-tolerant, load-balancing, efficient, extensible, and transparent.

But it's a lot easier to just say "FLEET."

Go with the QNX FLEET network.

Nothing runs like it.

To find out how your applications can thrive in the QNX environment, call **1-800-363-9001** (ext. 103).

ARCHITECTURE MAKES THE DIFFERENCE.™

QUANTUM SOFTWARE SYSTEMS LTD. 175 TERENCE MATTHEWS CRESCENT, KANATA, ONTARIO K2M 1W8 CANADA TEL: 613-591-0931 FAX: 613-591-3579

QNX IS A REGISTERED TRADEMARK AND FLEET IS A TRADEMARK OF QUANTUM SOFTWARE SYSTEMS LTD. © QUANTUM SOFTWARE SYSTEMS LTD. 1992

CIRCLE NO. 31

Real-time environments aimed at cost-sensitive 32-bit designs

Tom Williams, Senior Editor

There's a tacit assumption that any embedded application using a 32-bit microcontroller is high-end. But is a \$2,000 laser printer a high-end application when compared to a \$250,000 medical instrument—just because both use an Intel 960CA? The truth is, a growing number of 32-bit embedded CISC and RISC applications are extremely cost-sensitive because they address high-volume markets, such as intelligent cellular phones, fax machines and radios, while using as little memory and I/O as possible.


"Generally, applications are growing more complex," says Linda Thompson, director of product planning for Ready Systems (Sunnyvale, CA), "but a certain class of applications is growing more cost-sensitive." So the development environment—host/target communications and tool interface—needs to make minimum demands on the target hardware.

Developers of what we can call these low-end 32-bit systems don't want to build an Ethernet interface into their products if that interface won't be needed in the final production versions. They don't want to pile a lot of memory onto their boards if it's only going to support the cross-development process. And yet they need to have a multiuser environment in which different developers can access the target with different tools, because the complexity of the applications they're building continues to increase.

Unburdening target hardware

According to Tony Barbagallo, director of product marketing for Wind River Systems (Alameda, CA), the approach that's needed "is to move the equivalent level of functionality into the host-level debugger and have it communicate with a run-time agent on the target that's probably less than 20 k[bytes]."

Both Ready Systems and Wind River are introducing cross-development environments for low-end 32-bit systems. The two companies have adopted similar strategies, minimizing the burden on the target hardware. Instead of using networks and very small target agents to communicate with host-based tools, both have introduced systems that use serial links instead. Ready



"Spectra provides uniform tools and environment for all types of embedded software development," says Inbar Lasser-Raab, senior product marketing engineer of Ready Systems. "It satisfies the complete range of real-time applications from bare machines to RTOS-based systems."

Systems' new environment is called Spectra, and Wind River has dubbed its system MicroWorks. They differ in a number of ways, the most significant being that Ready's Spectra is intended as an open environment that will support third-party and proprietary real-time operating systems (RTOSs), while Wind River's MicroWorks remains bundled with its VxWorks development environment and its Wind real-time kernel.

Common to both environments is serial communications, which isn't as reliable as TCP/IP because it sacrifices error checking for low bandwidth and target overhead. In each case the host is only notified when

changes occur, reducing overall communications traffic.

Ready Systems' approach to host/target serial communications has been to establish a serial protocol called Xtrace. It's a user datagram protocol, part of TCP/IP's lowest level; it can run over a serial link or on networks. Xtrace communicates between a host-resident target manager and the Xtrace Daemon resident on the target system.

The Xtrace Daemon is independent of the target operating system, occupies less than 30 kbytes of target memory and is easily portable, even to custom hardware. "All it needs is one driver for a timer and

one driver for some kind of physical connection to bring it up," says Ready Systems' Thompson. "In most environments, all you need is a working driver and a working real-time OS to get a connection."

On the host side, tools communicate with the target manager, and so with the entire Ready Systems Spectra environment, via Spectra's ToolBuilder interface. This interface is an open, published set of about 70 calls that lets tool vendors adapt their products to Spectra. "We've checked with Microtec and Green Hills/Oasys," says Ready Systems' president James Ready, "and all of them have said the ToolBuilder interface is straightforward to move to."




In Wind River's MicroWorks, the VxGDB remote debugger has been enhanced so it communicates via a simple binary-encoded serial protocol that can go over RS-232. "Because we have complete and free access to the source code of our source-level debugger, we were able to modify it so that it can communicate with the very low resource requirement ROM monitor called VxMon," says Wind River's Barbagallo. VxMon is an external agent to the VxWorks real-time kernel. "Since all the symbolic information is now on the host," he adds, "we can also enhance the debugger to talk to emulators, because in some cases developers don't even have serial ports on their boards."

Tools on the host side of MicroWorks communicate directly with the

You Can't Debug Optimized Code Until You Have Optimized Code.

Only Green Hills Delivers Optimized Code.

↓ The level of optimization that should be free.




Dhrystone 2.1 Benchmark Results

These compilers were tested with a complete suite of industry standard benchmarks.
All tests were executed on a Motorola VME165 (25 MHz 68040 with caches in copy-back mode).
Green Hills consistently outperformed the competition.

Call for a complete benchmark report.


Green Hills
• SOFTWARE, INC. •
The Leader in Optimizing Compiler Technology


Oasis

1 Cranberry Hill, Lexington, Mass. 02173
(617) 862-2002 • Fax (617) 863-2633


Microtec is a registered trademark of Microtec Research, Inc. / Motorola is a registered trademark of Motorola, Inc. / C-68000 is a trademark of Green Hills Software, Inc. / Copyright 1992 Green Hills Software, Inc.

SOFTWARE AND DEVELOPMENT TOOLS

VxMon monitor over the serial link. A debugger or other host-based tool will have to be modified to talk to the VxMon interface which, like the ToolBuilder interface, will be a published specification.

Life before prototype

Ready's approach to host-based prototype development has been to port its Xtrace Daemon rather than its real-time kernel, VRTX, to a UNIX process. That process then becomes what Ready calls the virtual target. The virtual target provides the same operating system and I/O services that you'd have on a real target. Allowing prototyping at the C level, it isn't itself an instruction-set simulator, although you could conceivably put an instruction-set simulator in the virtual target. Host-based tools interact with the virtual target via the ToolBuilder interface in exactly the same way that they do



Ready Systems' Spectra environment supplies a common interface for development and debugging tools on the host system. Tools communicate with the real or a virtual target via the same interface. The shaded blocks identify the basic Spectra components.

when talking to the real target via the Xtrace Daemon.

Wind River, on the other hand, has "developed a full implementation of VxWorks that runs self-contained on a UNIX workstation," says

Barbagallo. "What the user sees is just what he or she would see if a target were hooked up."

This implementation of VxWorks, which is called VxSim, is basically a port to a Sun workstation, where the

Benchmarks at
107MHz!
Landmark V2.0 at 33MHz

Cyrix™ 486DLC

Passive Backplane - Single

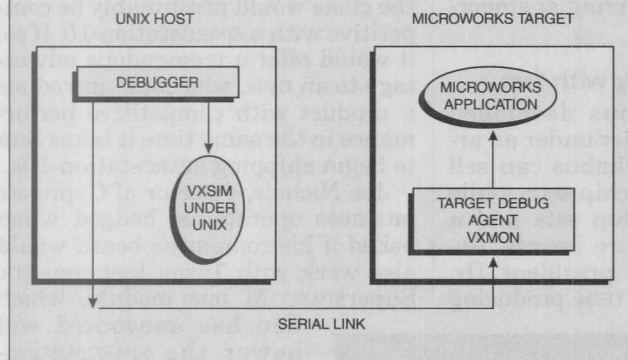
The photograph shows the Cyrix 486DLC Passive Backplane circuit board. Several components and features are labeled with callouts:

- 486DLC 33MHz "Engine Only"**: Points to the main CPU chip.
- CAT985DLC**: Points to a cache chip.
- Up to 32Mb of DRAM Onboard**: Points to the memory modules.
- IDE Interface**: Points to the IDE connector.
- SCSI Interface**: Points to the SCSI connector.
- Floppy Interface**: Points to the floppy disk connector.
- Up to 128Kb 2-Way Set Associative CACHE**: Points to the cache chip.
- High Performance Cyrix Cx486DLC**: Points to the main CPU chip.
- 83D87 Cyrix Math Co-Processor**: Points to the math co-processor chip.

Other visible components include an **auitek** chip, a **CAT990** chip, and a **VLSI FUTURE** chip.

Product names are trademarks or registered trademarks of their respective companies.

MicroWorks environment



software runs as a UNIX process using the SunOS lightweight process mechanism to manage context switching. "You're writing C source and debugging with VxGDB, whether you do it on the host or target," Barbagallo says. "When the hardware comes in, you just recompile and download."

Both Ready Systems' Spectra and

Wind River's MicroWorks let you do remote debugging over a network while simultaneously shielding the target from being connected directly to the network. A remote user, for example, can talk to MicroWorks on a machine to which a target is attached via the host's serial link, in this way debugging the target. The same holds true for Spectra. You can

The MicroWorks environment from Wind River Systems uses a UNIX process-based version of VxWorks as a host-based prototyping environment for code development. Tools communicate directly over a serial link with the target-resident VxMon, an external agent that is specific to the VxWorks RTOS. VxMon occupies only about 20 kbytes of target memory.

talk to the ToolBuilder interface on the target's host system.

One difference between the two environments is that Ready's Target Manager has a mechanism for setting up multiple channels so that multiple tools can access the same target simultaneously. Another is that the Xtrace Daemon operates independently of the target's real-time operating system.


Both Wind River and Ready Systems have correctly perceived that big things come in small packages. Reducing the final cost of product hardware will help them in the cut-throat race going on in the market in general and the consumer market in particular. ■

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

Ready Systems
(408) 736-2600 Circle 205
Wind River Systems
(510) 748-4100 Circle 206

33MHz Board Computer

Fully Integrated
Engine Only



| 386 FUNCTIONS | CAT990 Fully Integrated | | CAT985 Engine Only | |
|--|----------------------------|-------------|-----------------------|-------------|
| | without DLC | with DLC | without DLC | with DLC |
| 25/33MHz 386 | ✓ | | | |
| 33MHz Cyrix 486DLC | | ✓ | | ✓ |
| 25MHz 386 | | | ✓ | |
| Internal Cache | | ✓ | | ✓ |
| 64Kb-128Kb Optional Cache | ✓ | ✓ | | |
| Up to 1 Mb PROM/FLASH Disk | | | ✓ | ✓ |
| Up to 32Mb DRAM Onboard | ✓ | ✓ | ✓ | ✓ |
| 2 Serial Ports / 1 Parallel Port | ✓ | ✓ | | |
| SCSI/IDE/Floppy Interface | ✓ | ✓ | | |
| Noise Reduction Circuitry For FCC Class B | ✓ | ✓ | ✓ | ✓ |
| PS/2 Mouse Support | ✓ | ✓ | ✓ | ✓ |
| /AT-PS/2 Keyboard Support | ✓ | ✓ | ✓ | ✓ |
| Double-Sided Surface Mount Technology | ✓ | ✓ | | |
| Manufactured In-House (USA) | ✓ | ✓ | ✓ | ✓ |
| Landmark V2.0 at 33MHz | 50 | | | |
| Landmark V2.0 at 25MHz | | | 36 | |
| Landmark V2.0 at 33MHz with Cyrix 486DLC | | 107 | | 107 |

Diversified Technology®
An Ergon Co.

Call us toll free for orders and information.

1-800-443-2667

U.S.A. - (601) 856-4121 Fax (601) 856-2888
Outside U.S.A. - (201) 891-8718/Fax (201) 891-9629

Chip sets could unleash workstation clones

Don Tuite, Senior Editor

Will there be an explosion of workstation clones in the 1990s to match last decade's explosion of PC clones? Processor developers, such as Cypress and MIPS, are hoping for precisely that. But whether workstations relegate PCs to the fossil realm of Z80-based CPM machines or Intel dominates the market through the rest of the decade depends on a lot of factors, many of them software-related.

Setting software aside for the moment, however, one thing is clear. Unless a number of systems manufacturers who can fill the clone-maker niche for workstations step forward, Intel has little to worry about. The existing players (Digital, Hewlett-Packard, IBM, and Sun) aren't structured for the same kind of high-volume manufacturing and distribution as are the manufacturers of PC compatibles.

Of the existing workstation processor architectures, only Sun's and MIPS' are readily clonable open systems. On the Sun side, three announcements mark a new opportunity for systems OEMs who covet a piece of the SPARCstation market. One is from LSI Logic (Milpitas, CA), whose SPARKit chip set facilitates the manufacture of inexpensive SPARCstation-2 compatibles. The other two are from Cypress Semiconductor (San Jose, CA) and Nimbus Technology (Santa Clara, CA), who've announced separate packaging of the same chip set, each aimed at different-sized OEMs.

The Cypress/Nimbus chip set can handle dual SPARC processors, and could potentially let OEMs create their own versions of both the SPARCstation-2 and the SPARCstation-10. In a bid to entice the manufacturing-savvy Taiwanese and Korean PC clone-makers, both Cypress and LSI Logic are providing manufacturing kits that help OEMs reduce their sys-

tem-design non-recurring engineering (NRE) costs.

Performance parity with Sun

Cypress and Nimbus developed their products together under an arrangement where Nimbus can sell finished boards and chip sets, while Cypress may sell chip sets and a license to manufacture boards. According to Nimbus president Dr. Sanjeev Renjen, an OEM producing

the clone would presumably be competitive with a SPARCstation-10. If so, it would offer a tremendous advantage to an OEM, who could introduce a product with competitive performance in the same time it takes Sun to begin shipping SPARCstation-10s.

Joe Nichols, director of Cypress's business operations, hedged when asked if his company's board would also work with Texas Instruments' SuperSPARC M bus module, which

Sun has announced will power the SPARCstation-10. Nichols said that he didn't have complete knowledge of the TI architecture, but guessed that, "It probably would, although some minor operating system tweaks might be necessary."


When using existing Cypress processor modules, the design appears to be robust. Cypress has shown a single-processor prototype with system clock increased from 40 to 50 MHz simply by changing the crystal, but with Nichols noting that the faster interface was no longer in strict compliance with the 40-MHz M bus standard. With the faster clock, the demonstration unit calculated and displayed a Mandelbrot set sequence faster than a nearby SPARCstation-2.

Cheap to build

Forrest Lai, vice-president of workstation development for Tatung (San Jose, CA), says that before evaluating LSI Logic's SPARKit, "We didn't think it was possible to develop a system selling for under \$5,000." Now, he expects that such SPARCstation-2-compatible systems can be "highly competitive."

LSI Logic calls SPARKit a uniprocessing solution, although the company acknowledges that the hardware supports M bus Level 2. Level 2 is required for multiprocessing as it's performed in the SPARCstation-10.

Cypress, with its separate integer unit, floating-point unit and multiprocessing cache controller/MMU, fa-



James Warford, Nimbus Technology's manager of systems integration, says, "Using 'known-good' M bus processor modules is an advantage in wringing out a system. M bus timing is well-defined, and troubleshooting is relatively straightforward." Here Warford uses a header between a motherboard and processor board to put a logic analyzer on the M bus.


100 to 200 systems per month would find it more economical to deal with his company, while manufacturers who anticipate larger volumes would be better served by Cypress.

When provided with a power supply and peripherals, a Cypress SPARCset or a Nimbus NIM6000M chip set with a motherboard mates with Cypress M bus modules to produce a complete SPARCstation compatible. With a single-processor M bus module, the workstation is roughly equivalent to a SPARCstation-2. With a functional dual-processor Cypress HyperSPARC module,

FPGA Speed Limit Now Upped To 150 MHz.

That's AT&T "Customerizing."

AT&T 3000 Series FPGA Performance Trends



Introducing AT&T's 150 MHz FPGA, equipped with the highest available speed to support the fastest microprocessors. Ready for design-in right now, it offers 2000 to 9000 gate densities, so you won't have to compromise system speed with slow logic as your microprocessor gets faster.

AT&T's 150 MHz FPGA also provides you with the lowest possible cost to support your bottom line. And because it's designed for the most

AT&T 3000 Series FPGAs 150 MHz Toggle Rates

popular FPGA architecture, you can design without worrying about the future.

Our 150 MHz FPGA is just another step

in AT&T's continuing dedication to offer even faster FPGAs priced even lower. In 1993, we'll deliver

ATT3000-Series FPGAs that can handle 66 MHz microprocessor speeds.

Meanwhile, take advantage of the substantial savings now being offered on AT&T's ATT3000-Series 150 MHz devices. They're available from AT&T Microelectronics or your authorized AT&T distributor.

For a free 150 MHz Speed File to update your XACT* Development System, plus our FPGA Information Package, complete and mail the attached card, or call AT&T Microelectronics at 1 800 372-2447, ext. 905.



AT&T
Microelectronics

*XACT is a trademark of Xilinx, Inc.

CIRCLE NO. 34

INTEGRATED CIRCUITS

vors M bus modules that plug into the motherboard, and Sun's MP690 servers use pairs of two-processor Cypress modules. Following this design strategy, SPARCset and the NIM6000M provide M bus headers on their motherboards, and the chip sets support M bus Level 2.

Although both the LSI Logic and the Cypress/Nimbus chip sets comprise six devices, there are numerous differences that reflect the different price/performance targets of the two offerings. Emphasizing its low-cost approach, LSI Logic uses an SBUS-based SVGA graphics controller. Cypress's controller resides on M bus, which makes it faster, according to Nichols. It can handle 1152 × 900 pixel resolution and true color.

The integrated LSI Logic MMU/cache/cache tag chip lives on the motherboard. Cypress cache/MMU chips and separate cache SRAMs live on the processor daughterboard. Among other differences, Cypress

has split its DRAM controller into a pair of chips, so that each handles 32 bits of the 64-bit M bus datapath. Two chips can address 128 Mbytes of DRAM. LSI Logic went with a single chip and addresses 8 Mbytes.

LSI Logic's I/O is the same as SPARCstation-2. Cypress and Nimbus took a different approach. Nichols says that when the program was just getting under way, the current explosion of SBUS-based peripheral cards hadn't begun, and the designers' approach was to take advantage of the multitude of peripherals provided for the PC. Consequently, two chips in the SPARC/NIM6000M chip set provide an interface between M bus and the 386SX bus. This is still a low-cost avenue for providing keyboard and mouse, and two SBUS slots allow for more sophisticated peripherals.

Fast turn, low NRE

Taking a leaf from the PC chip set makers' notebook, both Cypress and

LSI Logic can now offer OEMs complete packages to get them past the hurdle of designing a 40-MHz M bus board, so they can move on to developing manufacturing economies and other product differentiators. These latter include drawings, Gerber tapes, copyable boot ROMs, diagnostics, and development copies of the OS.

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

| | |
|------------------------------|------------|
| Cypress Semiconductor | |
| (408) 943-2600 | Circle 207 |
| LSI Logic | |
| (408) 433-8000 | Circle 208 |
| Nimbus Technology | |
| (408) 727-5445 | Circle 209 |

2ND ANNUAL ANALOG & MIXED-SIGNAL DESIGN CONFERENCE



100% DEDICATED TO YOUR NEEDS

If you are an analog design engineer, a digital designer who must address the analog implications of high speed design, or an engineering manager, you should attend this conference.

Choose from the widest array of lectures, workshops, panel discussions, tutorials, and vendor information available at a single event. You'll learn about analog and mixed-signal simulators, different design techniques, analog implications of high-speed systems, and much more such as:

- How to bring mixed-signal designs to market quickly
- New innovative solutions to tough design problems
- The capabilities and limitations of analog and mixed analog/digital ICs, ASICs and CAE/CAD tools
- Secrets to boosting your productivity
- Cost effective ways to implement high-speed digital, mixed-signal, and analog designs



ANALOG & MIXED SIGNAL

1010100101001010010001001
000110100100010101110100
1000101010100010010101
01000101010101010101011
010001010100010101001010
1100110001010001101100001
0110001010011001011000110


DESIGN CONFERENCE

JOIN YOUR COLLEAGUES THIS OCTOBER

Don't miss this opportunity to strengthen your skills and increase your designing power.

For more information call Tom Loftus at 415-905-2354 or fax 415-905-2220.

The competition will call us ruthless. You can call us at 1-800-234-4VME.



It's enough to make other VME board builders call us names. Or call it quits. A new 38 MIPS* VME single board computer based on the 88100 RISC microprocessor. Or a new 26 MIPS* VME board based on the 68040 CISC microprocessor.

Both are built by Motorola and offered at \$3,995 each. That's just \$105/MIPS for the RISC board, which compares nicely with the \$1,000/MIPS you've been asked

to pay for somebody else's board. And it's just \$154/MIPS for the CISC board.

The MVME187 (RISC) and MVME167 (CISC) boards employ VME D64 architecture. And both come with four 32-bit timers.

For a free color brochure, call the 800 number above. And see why the competition undoubtedly wishes we'd call the whole thing off.



MOTOROLA
Computer Group

Motorola and the M are registered trademarks of Motorola, Inc. © 1992 Motorola, Inc. All rights reserved. *MIPS rating based on Dhrystone 1.1 test results where 1757 Dhrystones/s is 1 MIPS (VAX™ 11/780). MIPS performance is based on the Diab 2.36E compiler.

CIRCLE NO. 36

Specialty PROMs count bursts, implement state machines

Don Tuite, Senior Editor

Memory chips were once distinguished exclusively by array size and access time. This situation has started to change as manufacturers introduce processor-specific static RAMs (SRAMs). SRAMs that target the 80486 or 68040, such as Motorola's 32 kbits \times 9 MCM62486A and MCM62940A, provide not only glueless interfaces to their respective processors but also integral counters that deal with burst sequences.

Now there are specialty PROMs as well, and they offer more than just a clean interface to processors. Cypress Semiconductors (San Jose, CA), for example, has just introduced both a latched and registered PROM with a counter and an unusual PROM with user-programmable input and output registers that can be used to create fast state machines.

General purpose, no glue

Although Cypress bills its 16 kbits \times 16 CY7C270 as a processor-specific PROM, the chip's real strength is the universality of its glueless interface. If you're using a RISC processor that asserts addresses at the start of a clock cycle and then drops them, the PROM will register them. If you're using an 80X86 or a 680X0, where addresses are available before the clock leading edge, you can latch them early.

For 486 burst accesses, you configure the burst counter to use an integral look-up table. For most other processors, you configure the counter in its 2-bit mode, and for the 29000, you configure it in 8-bit mode. (The counter can also be set for four bits.) Processors that suspend bursts can signal the PROM counter to start and stop as necessary.

Because there's no glue logic between the processor and program memory, you really can use these PROMs to eliminate shadow RAM in high-speed embedded applications. The fastest access time for burst reads is 14 ns; for single reads, it's 28 ns. This is for the -20 speed bin of the chip. To deal with its burst

capability, the speed designator in the CY7C270 nomenclature refers to the minimum clock period rather than to the PROM's maximum access time.

Cypress says that the $\times 16$ organization reduces parts count in 16- and 32-bit processor designs. Because you can program the polarity of the three chip-select inputs, you

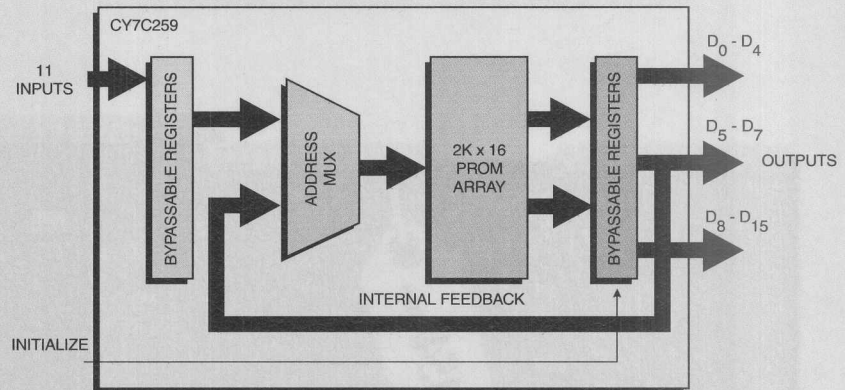
can provide up to eight banks of 16 kbits \times 16 PROM.

No logic minimization

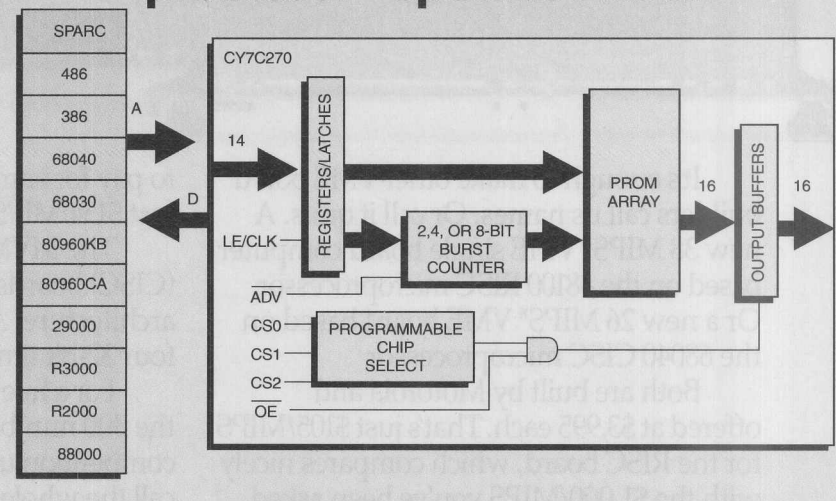
Although they bear architectural similarities to the processor-specific CY7C270, Cypress' other new specialty PROMs, the CY7C258 and CY7C259, aim at an unusual and more specialized niche. They're designed for a certain class of state machines. (The difference between the two chips is the number of external outputs—8 or 16.)

Cypress says the advantages of its state-machine PROMs (SMPs) are best

Cypress 83-MHz state machine PROM



Cypress programmable processor-specific PROM



By integrating registers, feedback and counters with fast PROM cores, Cypress has created two high-speed specialty PROMs. One (A), for complex state machines with relatively few inputs (11 max), can run at 83 MHz. The other (B), for instruction storage, can service virtually any of the fastest RISC and CISC processors.



Now You Can See True Colors Without Getting Soaked.

New full-featured RAMDACs from Brooktree deliver 24-bit true color for cost-conscious PC designers. Introducing four new RAMDACs that span the spectrum of PC applications, from 640x480 VGA systems to 1280x1024 workstation-quality graphics. They've got the right features, the right prices and are available right now.

True Color in a VGA Environment

Unplug Sierra. Plug in our totally compatible Bt481 or Bt482. You'll get 24-bit performance at 16-bit prices.

These new RAMDACs support 15-bit TARGA, 16-bit 5:6:5 and 24-bit true-color formats. They even allow you to switch between VGA and true color on a pixel-by-pixel basis.

Choose the Bt481 if you prefer an external

hardware cursor. Or pick the Bt482 for its on-board 32x32x2 cursor — ideal for faster windowing environments.

Workstation Graphics at PC Prices

Introducing the Bt484 and Bt485, our newest mouth-watering RAMDACs. They combine true color with higher resolutions for higher performance graphics subsystems. And they're economical, too.

Bt485 operates at up to 135 MHz to drive pseudo color to 1280x1024 resolutions and gamma correct true color to 1024x768. It has a 64x64x2 cursor and all the serialization and timing to directly interface to a VRAM frame buffer.

Bt484 provides maximum flexibility with its programmable pixel port to provide 256

to 16.8 million colors, on-board 32x32x2 cursor and supports both interlaced and non-interlaced monitors.

| Bt | Resolution | | | |
|-------------------|-------------|---------|----------|---------------------|
| | 640x480 | 800x600 | 1024x768 | 1280x1024 |
| Colors | Bt481/Bt482 | | | Bt485 |
| 16.8M 64K 256 | Bt484/Bt485 | | | Workstation RAMDACs |

Call Brooktree at **1-800-VIDEO-IC** for technical details and pricing today.

Brooktree Corp., 9950 Barnes Canyon Road, San Diego, CA 92121, (619) 452-7580, FAX (619) 597-0673.

Brooktree®

INTEGRATED CIRCUITS

seen in state machines that are fairly big, since the SMP can accommodate 2,048 unique states, but that are modest in terms of inputs. If you do use a PROM rather than a PLD for these state machines, you don't have to perform any state minimization. That is, if the number of

state feedbacks plus the number of inputs is equal to or less than eleven, which is the number of address pins, the design will fit.

Of course, designers have been putting state machines into PROMs for years, using external logic to register and feed back outputs. Typical

PROM access times, added to the propagation delays of external logic, however, have made these designs rather slow. In contrast, Cypress claims that signals can be clocked at 83 MHz in the CY7C258/9.

Typical of state machines that could be implemented in a PROM are pattern generators and microcode sequencers. With a pattern programmed into the PROM core, you can use several additional bits in each word to store the addresses of the next PROM location. Then, to step from one state to another, you feed back these bits as inputs, so that a new output pattern is selected on each clock cycle.

If you're using the device as a microcode sequencer, the PROM can store code divided into blocks. Depending on the address presented, the state machine steps sequentially through the corresponding block.

Internal multiplexers

Cypress uses eleven 2-input multiplexers between the input registers and the PROM cores. Each multiplexer can be configured to pass either the input address or an internal feedback signal to the PROM array.

In both devices, the memory array is organized as 2 kbits \times 16. The 28-pin CY7C258 supports eight direct outputs. The other eight outputs from the PROM array are internally fed back to on-chip input multiplexers. The 44-pin CY7C259 has 16 outputs.

In both chips, there are three groups of outputs. In both, D0-D4 go straight from the array to the output pins without feedback. Also, D5-D7 in the two SMPs are routed to both the output pins and the address multiplexers. On the CY7C258, D8-D15 are used exclusively for feedback to the address multiplexers and have no corresponding pin connections. On the CY7C259, D8-D15 are routed to output pins.

Both state-machine PROMs have input and output registers. As is the case with the processor-specific PROM, input registers are useful for capturing signals that appear for only a short time. Output registers are used to hold state information or output data, or to meet the data setup/hold requirements of other devices.

Each bit in each register can be individually configured so that it

Finding the best operating system for your REAL-TIME application is a matter of simple arithmetic:

Start with proven operating system performance from Pdos, the leader in real-time response as validated by independent benchmark tests.*

+ Add development in the Pdos Full Spectrum™ development kit, which provides all the powerful and debugging tools you'll need to generate cross or native code that can be shared by a variety of platforms like MS-DOS, SUN, HP and VAX.

+ Connect it to a suite of networking and backplane facilities for flexible development and run-time communications.

+ Include "lifesaving" support from Eyring's real-time experts to help you keep your project afloat and on schedule.

+ Figure in the ease of use and effortless hardware integration Pdos offers to increase development productivity and minimize project costs.

The ultimate operating system and development tools for your real-time application = PDOS

Call 1-800-YES-PDOS

EYRING
CORPORATION

*Independent benchmark test results available upon request

CIRCLE NO. 38

Graphic memory that lets your imagination run wild.

| Memory Configuration | Features | Part Number | Access Time (ns) | Packaging Options | Availability |
|--------------------------------|--|------------------|------------------|-------------------|--------------|
| Triple Port DRAMs | | | | | |
| 256K x 4 | Fast Page Mode/Block Write | MT43C4257/8 | 80,100 | SOJ | Now |
| 128K x 8 | Fast Page Mode/Block Write | MT43C8128/9 | 80,100 | PLCC | Now |
| Dual Port DRAMs (VRAMs) | | | | | |
| 256K x 4 | Fast Page Mode/Block Write | MT42C4256* | 70,80,100 | ZIP, SOJ | Now |
| 128K x 8 | Fast Page Mode/Block Write | MT42C8128* | 70,80,100 | SOJ | Now |
| 256K x 8 | Extended Data Out/Block Write/Programmable Split | MT42C8256 | 70,80 | SOJ, TSOP | Now |
| 256K x 8 | Fast Page Mode/Block Write | MT42C8255 | 70,80 | SOJ, TSOP | Now |
| 256K x 8 | Fast Page Mode/Block Write/Dual Write Enable | MT42C8254 | 70,80 | SOJ, TSOP | Now |
| Specialty DRAMs | | | | | |
| 256K x 16 | Fast Page Mode | MT4C16256/7/8/9* | 70,80 | ZIP, SOJ, TSOP | 4Q92 |

*Low power versions also available.

Windows™ is a trademark of Microsoft Corporation.

MICRON
SEMICONDUCTOR, INC.

2805 E. Columbia Rd., Boise, ID 83706 (208) 368-3900
Customer Comment Lines: U.S. 800-932-4992; Intl: 01-208-368-3410

©1992, Micron Semiconductor, Inc.

CIRCLE NO. 39

Technology
that works for you.

will capture its input at the rising edge of the clock, or so that it's bypassed. If you bypass both sets of registers, the SMP becomes an asynchronous PROM. In that case, regardless of the clock, the PROM outputs will respond to any change at the inputs after the propagation delay.

The PROMs also have an asynchronous output-enable (OEB) and a synchronous chip-select. Like the registers in the datapath, the chip-select register can be individually bypassed. Other uses

Although the CY7C258/9 are aimed at state-machine applica-

tions, they may also be useful for program storage in very high-performance embedded systems using pipelined processors. The SMP is the only PROM with both input and output registers, making it a pipelined memory. It's fast enough to replace shadow RAM in high-speed applications. In fact, its clock-to-output time of only 9 ns makes it possible for the SMP to provide one instruction per clock cycle to pipelined processors running at clock rates as high as 83 MHz.

The first step in adding complexity to memory chips was taken by manufacturers when they did a glue sweep for each popular processor. This resulted in processor-specific SRAMs, such as Motorola's 32 kbits + 9 devices. Cypress designers have now taken the next logical step by adding more general-purpose logic to plain memory.

What's particularly interesting is that the memory they chose was the PROM. The result in one case is a fast PROM that's likely to be popular in PostScript printers and similar embedded applications. In the other case, the result is really a new class of part. Most electrical engineers have considered the PROM as a state machine only in passing, and not very practical because of the speed penalty imposed by the need for external feedback. With this penalty removed, PROM-based state machines may soon become more common in a variety of applications. ■

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

Cypress Semiconductor
(408) 943-2600 Circle 210

OOPS!

The Technology Focus article entitled "DSP boards reach performance highs" (*Computer Design*, August 1992 issue) included an incorrect telephone number, (301) 572-7801, for Ixthos in Silver Springs, MD. The correct number is (301) 572-6700.



The Top 6 Reasons You Should Call MIZAR Before You Start Your Next Embedded Design.

1 You're smart. You want the solution that is best for your design.

2 You want to get your product to market as quickly as possible. Mizar's VMEbus solutions give you a proven hardware starting point to build on. And Mizar's variety of real-time software solutions enable you to get your system up and running in record time.

3 You want to minimize design risk for your product. With Mips-based R4000 or R3000 CPUs, SPARC on VME designs and a variety of 68K ('010 to '040, including '340) single board computers, Mizar can supply the technology that you need. Whatever processor you have chosen, Mizar has the solution.

4 You need maximum design flexibility. With Mizar's VME MXbus you can add custom I/O to a proven VME design and create a custom CPU card or a smart I/O card. Even if you need a custom product, Mizar has the solution.

5 You can get everything you need from Mizar including computer I/O such as SCSI, Ethernet or graphics controllers. And Mizar has all manner of industrial I/O as well, including digital and analog I/O, machine vision and motion control products. So if you need I/O, Mizar has the solution.

6 Mizar's solutions will work in your system to make you and your product star performers. Whatever your reason, call Mizar today.



1-800-635-0200

MIZAR

Your Source for Solutions

2410 Luna Road
Carrollton, TX 75006
(214) 277-4600 (214) 277-4666 FAX

Mizar is a registered trademark of Mizar Digital Systems, Inc.
Other names are trademarks of their respective manufacturers. © 1992 Mizar Digital Systems, Inc.

CIRCLE NO. 40

Leaders in mixed-signal design to meet in San Francisco

Stephan Ohr, Contributing Editor

High-level description languages for analog, mixed-signal simulation models and techniques, mixed-signal scan test, and new applications for mixed-signal devices will head the agenda when the 2nd annual Analog and Mixed-Signal Design Conference convenes at the San Francisco Airport Hyatt Regency, October 28-30. Jointly sponsored by Miller-Freeman and *Computer Design*, this year's conference is covering new ground in tracking the progress of AHDL, evaluating analog and mixed-signal simulators and ensuring the effectiveness of mixed-signal test.

Who's coming to dinner?

All of the big guns in analog and mixed signal IC technology are providing speakers: Analog Devices, AT&T Microelectronics, Burr-Brown, Fujitsu Microelectronics, GEC Plessey, Gould AMI, Harris Semiconductor, National Semiconductor, NCR Microelectronics, Philips/Signetics, SGS-Thomson, and Texas Instruments. All the mixed-signal tool makers are represented with speakers or panelists: Analog, Cadence Design Systems, Dazix, Mentor Graphics, Meta-Software, MicroSim, Racal-Redac, and Viewlogic Systems. Even manufacturers whose expertise is specialty analog rather than mixed-signal—Elantec and Linear Technology, for example—will provide top-flight speakers this year.

While the selection of vendors and the design of mixed-signal ASICs dominated the presentations of IC manufacturers at last year's conference, this year's event will capture the new variety and flavors of mixed-signal applications. Computer graphic RAMDACs, disk drive read channels, digital audio playback converters, wireless transceivers, power op amps, switched-capacitor circuits, and DSP filters will be among the circuits presented. To be sure, this year's conference will also discuss technology choices for mixed-signal ASICs—Bi-

CMOS, CMOS, smart power processes—even ultrahigh-speed bonded-wafer processes. But the orientation of the ASIC presentations has shifted away from introductory material toward more sophisticated applications and problem solving.

One of the best presentations on process technology, for example, will be given by Steve Moore of GEC Plessey Semiconductors, who will describe partitioning alternatives for disk drive read channels. With disk drive form factors shrinking, there's pressure on semiconductor makers to build a single-chip disk drive. But not everything can be done in CMOS

or even BiCMOS. If Moore does his job well, he'll present a disk drive circuit, and discuss some of the ways it can be partitioned to utilize process technologies efficiently.

AT&T Bell Lab's Lauren Brust will point out some of the consequences the move to 3-V logic will have on analog and mixed-signal peripheral components. (Her presentation replaces Gus Richard's on the previously distributed schedule.) Texas Instruments' Lisa Schartz, similarly, will show how to build a rotary digital encoder from an optical sensor/amplifier cell. Fujitsu's Ron Gadway will demonstrate how wireless networks can be constructed from custom RF components.

Analog Devices will have a number of presentations on sigma-delta technology. TI's Daryl Sartain will show the application of sigma-delta

SGS-Thomson's Dauvin to keynote Analog & Mixed-Signal Conference

Jean-Philippe Dauvin, director of corporate market research for SGS-Thomson Microelectronics and chairman for European operations of WSTS (World Semiconductor Trade Statistics) will kick off the second annual Analog & Mixed-Signal Design Conference on Wednesday evening, October 28, with a keynote address on the current state of mixed-signal IC and ASIC technology and the future areas of growth for mixed-signal chips and systems.

"Dr. Dauvin has amassed an enormous amount of data on the mixed-signal market," said John Miklosz, editor-in-chief of *Computer Design* and co-chairman of the Conference, "and can provide insights based on sound analysis, not on speculation and wishful thinking. If anyone has been wondering about where the technology and markets are headed—who'll be the winners and the losers, so to speak—Dr. Dauvin can provide some answers."

Dauvin received a doctorate in economic sciences from Paris University's Institut d'Administration des Entreprises in 1966. A year later, he joined BIPE, where he created the Department of Electronics and carried out the first strategic studies on the electronic industry for the Commissariat au Plan of France and for the European Economic Community. He left BIPE in 1982 and joined the Thomson Group's strategic planning organization as chief economist. In this position, Dr. Dauvin was deeply involved in the rebirth of the Group itself, currently ranked number 10 in the world. In June 1987, following the merger between Thomson Semiconducteurs and SGS Microelectronics, he became director of corporate market research of the newly formed SGS-Thomson Microelectronics Group.

In addition to his strategic planning responsibilities at SGS-Thomson Microelectronics, Dr. Dauvin is chairman for European operations of WSTS and vice-president, worldwide. Based in San Jose, CA, WSTS operates as an independent, non-profit organization whose objective is to provide its members with detailed statistical data and accurate forecasts about the semiconductor market.



INTEGRATED CIRCUITS

to digital audio circuits. Philips/Sigmetrics' Enjeti Murthi will be the first to do a presentation on digital disk drive read channel implementation. There will be two presentations on DSP algorithms and IC filter development. One will come from the Warren, NJ startup, Star Semiconductor, whose DSP development tools are becoming famous for their ease of use.

Focus on design and simulation

One special advantage of the conference will be its focus on mixed-signal design tools and simulation issues. There will be three separate panel discussions covering the progress toward an analog hardware description language (AHDL), suggestions for benchmarking mixed-signal simulators and the use of simulation backplanes. ("If the typical designer uses more than one simulator," asks Hal Alles of Mentor Graphics, "what is the best way of tying them together?") Also featured are separate presentations on scan test techniques for analog boards and ASICs.

While the standardization of an AHDL may still be three to five years away, IEEE committee members and DARPA-funded agencies are already at work defining the necessary elements of an AHDL. The conference on analog and mixed-signal design is one of the first forums to bring AHDL participants together in one place. A Friday afternoon panel, chaired by *Computer Design*, will hear presentations from David Smith of Analogy, whose MAST modeling language is being evaluated as a foundation for both an AHDL and the DARPA-funded MIMIC (millimeter and microwave IC) hardware description language (MHDL); Dave Barton of Intermetrics, who has the DARPA contract to develop an MHDL that will be useful for "low-frequency analog;" Mark Brown of CLSI, who chairs the IEEE subpar committee charged with developing analog extensions to VHDL; and Rajeev Mandhavan of Cadence Design Systems, who is lobbying for an analog high-level language that capitalizes on the popularity of Spice. Other participants include Dr. Ed Cheng of Mentor Graphics and Peter Denyer of AnacAD, which offer high-level language alternatives to Spice.

The panel entitled "Guidelines for Benchmarking Mixed-Signal Simu-

lators" will take on some of the issues raised by last year's panel, "The BCTM Challenge," a simulation exercise discussed at the 1990 Bipolar Circuits and Technology Meeting. The BCTM presentations and subsequent arguments revealed that Spice, Saber and other modelers will not allow an easy "apples-to-apples" comparison. What does it mean when one simulation run took 9 min and 13 s on a VAX, another took 6 min and 56 s on an HP/Apollo, and a third took 3 min and 29 s on a Sun SPARC—but three days to tweak the netlist? In fact, a major research survey on mixed-signal recently completed by *Computer Design* found that users overwhelmingly demand simulation accuracy rather than speed. The panel session on benchmarking simulators will attempt to find consensus among simulation tool vendors rather than controversy—agreements as to what constitutes a good benchmark and why. Participants in this session include Mark Chadwick of Analogy, Dündar Duşlülöğ of Cadence, Graham Bell of MicroSim, and others. Participants and observers are encouraged to bring in circuits that may serve as industry-standard benchmarks.

High interest in test

Interest in mixed-signal test and design-for-test (DFT) is running high, gauged by the number of presentations offered at this year's conference (no less than four). Steve Dolens of IMP and Mani Soma of the University of Washington—co-chairs of IEEE 1149.4, the committee charged with assembling a scan test standard for mixed-signal testing—will do a presentation on the development of the standard. This will be one of the first public presentations on 1149.4 outside of the International Test Conference (and assorted DFT meetings). The presentation will be supplemented by one from Richard Hulse of Gould AMI Semiconductors, which participates in the 1149.4 committee and has begun developing a scan test methodology for mixed-signal ICs. Another 1149.4 participant, Stephen Bateman of Cadence, will do a presentation on the use of tester models with analog and mixed-signal simulators, a process guaranteed to improve DFT. Luke Hsieh (pronounced "Shay") of LTX, similarly, will do a

presentation on tester models that can be used in the design environment.

Back to basics

Among the most popular sessions last year were those that offered basic instruction on analog thought processes to digital designers. One repeat from last year will be Kerry Lacanette's half-day tutorial on monolithic A/D converters. (Lacanette, as some of you may be aware, was a Bob Pease protégé at National Semiconductor.) Another repeat from last year will be Bob Leonard's talk on frequency domain specifications—FFTs, harmonics, signal-to-noise ratios—an absolute must for anyone exploring the uses of DSP. Though this will be presented as a tightly-focused one-hour lecture, Bob's presentation is based on a popular 3-hour seminar that he gives for Dadel. A half-day tutorial, "Analog Filter Design," will be given by Geert De Veirman and Richard Contreras of Silicon Systems. Both Drs. De Veirman and Contreras are well-published in the field of filter topologies.

Also popular last year were the tutorials "The Basics of Spice" and "An Introduction to Transmission-Line Analysis." This year, the Spice tutorial will be given by Dr. Jeffrey Deutsch of Deutsch Research, who contributed to the development of the MOS "relaxation algorithms" for Spice at Berkeley. Having designed Spice simulators for both the Macintosh and the IBM PC, Dr. Deutsch remains one of the most accessible authorities on Spice. The transmission-line tutorial will be given by Dr. Al Wexler of Quantic Labs, and will be based on Quantic's popular half-day seminar.

Other talks on transmission-line analysis—an increasingly popular subject, especially for high-speed digital designers—will be given by Jon Powell of Quad Design and Dr. Paul Wang of Contec Microelectronics.

To receive a copy of the Analog & Mixed-Signal Design Conference program and registration information, call Patti Kenney at 1-508-392-2124. Remember, the Conference starts on Wednesday, October 28.


ANALOG & MIXED-SIGNAL DESIGN CONFERENCE PROGRAM-AT-A-GLANCE

9AM 10:15 11:30 12:30 1:30 2 3:30 4:30

WEDNESDAY

| | | | | | | | |
|---|--|--|--|---|---|--|--|
| 101 The Basics of Spice Modeling | | | | 131 Guidelines for Benchmarking Mixed-Signal Simulators | | | |
| 102 Analog Filter Design | | | | | 132 Front-end Components for Wireless Systems | | 142 Current Conveyors: A New Look at Some Older Applications |
| 103 Making the Transition to Digital Filtering Techniques | | 113 Trends in Operating Voltages for Mixed-Signal ICs | | | 133 Choosing PC Graphics Architectures by Application | | 143 Real-World Applications for Sigma-Delta Converters |
| 104 Mixed-Signal Tester Models | | 114 Technology and Partitioning Alternatives for Disk Data Paths and Servo/Spindle Control | | | 124 Concurrent Design-to-Test Methods for Mixed-Signal IC Development | | 144 Issues in Mixed-Signal CMOS ASIC Design |
| | | | | | 135 Behavioral Modeling of Logic Gates | | 145 DSP Design and Simulation |
| | | | | EXHIBITS OPEN 12:30-7:30 | | | |

Introducing the
real-time
development system
that won't
drive you to drink.



REDEFINING
REAL-TIME

Thanks to Spectra, you no longer need to rely on caffeine or those gooey simple carbohydrates to get your job done. That's because Spectra was designed to speed up software development. Not your heart rate.

Spectra is a bona fide breakthrough in real-time technology. With cross-development features that have never existed before. The Spectra tools work for all real-time applications. From bare machines to complex systems. From device drivers to multithreaded applications. And from microcontrollers to high-end microprocessors. So now, even lightweight targets can have heavyweight tools.

What's more, Spectra's Virtual Target lets you debug your application on the host before your hardware is ready. All of which makes Spectra ideal for custom designs.

And because the Spectra system is truly open, it supports our tools, your tools and third party tools. Finally, Spectra offers a 100% VRTX compatible, scalable run-time that can be tailored to your application.

So if you're ready to kick the caffeine habit, call 800-321-2636. Now, if we could just figure out a way to get you to lay off the Twinkies.

**READY
SYSTEMS**

© 1992 Ready Systems. Spectra, VRTX and Virtual Target are trademarks of Ready Systems. ® Jolt is a registered trademark of Jolt Company, Inc.

CIRCLE NO. 42

Performance analysis spots hardware/software bottlenecks

If your system works, but doesn't perform as expected, try performance analysis tools. They can provide a high-level view of system operation, zoom in on problems and help you avoid performance bottlenecks in the design stage.

Tom Williams
Senior Editor


It works. Well, sort of. So, how well does it work? Does it always work according to spec, or are there times—perhaps critical times—when something falls out of spec? How well are the hardware resources being used? Is hardware design contributing to software bottlenecks?

These are questions many designers have to ask themselves about their systems. To find the answers, you need tools that can look at the way a system performs, not just by analyzing registers and bytes, but by looking at the system's overall architecture and what the software is doing at the functional and algorithmic levels. You've got to know where the CPU is spending what portion of its time, how long critical routines are taking to execute and, in the case of multiprocessor systems, how well the load is balanced among CPUs. You need performance analysis.

Analysis of how well a system is performing can be a big help at integration time, when the software is married to the final target hardware. But analysis can also be a valuable asset at the design and simulation stages, when you're deciding just what the hardware execution environment should be. Maybe you don't need as fast a CPU as you thought. Maybe three \$20 CPUs will be more reliable than a single, high-powered \$250 CPU.

Tools that help at the integration stage are available to run on workstations, logic analyzers, emulators, and special analysis systems. Tools for estimating system performance at the design stage are a little farther behind, but are becoming available as CASE technology matures.

According to Phil Berger, a software engineer in the test and measurement



Tektronix's Phil Berger says there are trade-offs between using statistical and real-time performance analysis. Statistical analysis samples execution and could miss some critical events. Real-time analysis will catch every event, he says, "but with real-time you don't have the ability to prefilter certain cycles that you aren't interested in."

TECHNOLOGY FOCUS: PERFORMANCE ANALYSIS

division of Tektronix (Beaverton, OR), performance analysis from an external monitoring device, such as a logic analyzer or an emulator, can be classified in one of two categories: statistical or real-time performance analysis.

Limits of statistical analysis

"Statistical performance analysis isn't able to capture 100 percent of the activity on your system," says Berger, "because it has to go off and do something with the data it acquires." Real-time performance

mance analysis would probably meet all your needs."

Jay Maggard, director of marketing for Applied Microsystems (Redmond, WA), offers a caution, however, about statistical analysis. "If the interval between samples matches a period within the program, it can miss some things altogether," he says. If the event system is set up to sample every 1,000 bus cycles, and there's a loop in the program of about the same length, some small routine executed from that loop could appear to consume many more

behind it."

If a system absolutely requires that an interrupt be serviced within a specific time, you want to find out if there's ever a case where the process exceeds that time. To find out, the logic analyzer or emulator is set up to monitor a set of address ranges defining the function of interest. It measures the time of the function from the beginning event to the end. The best-case, worst-case and average times can then be displayed as a histogram. If any percentage of the repetitions of that function fall outside the requirements, it's an indication you need to go in and investigate why.

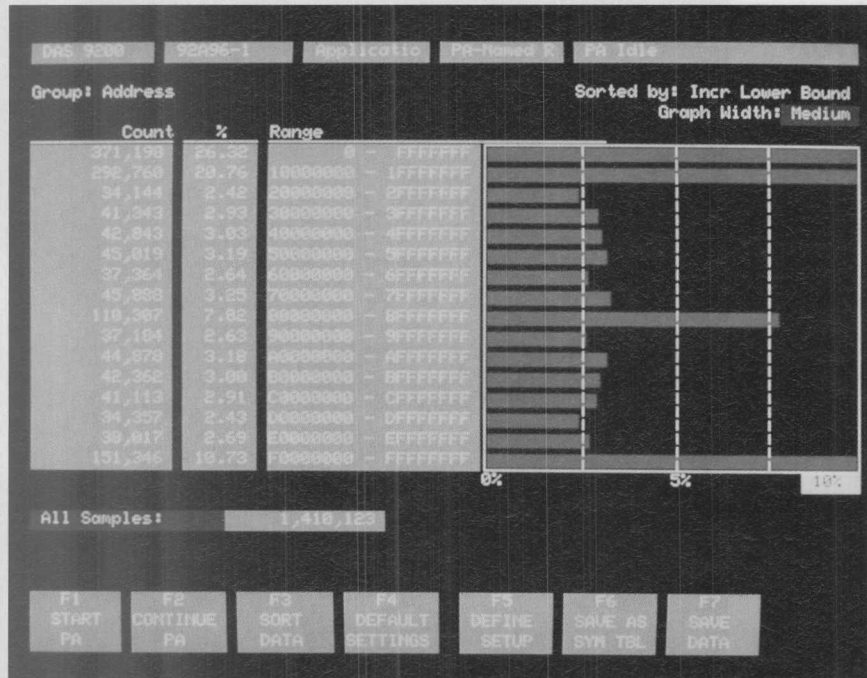
Analysis options

Tektronix offers performance analysis options for both its logic analyzer and digital analysis system (DAS) products. "The DAS is a true statistical analysis tool," says Berger. "It captures an entire buffer, processes that data and captures another buffer over and over." DAS 9200 systems work via a network and display their data in X Windows on a Sun workstation. With PA92 performance analysis software, you can specify ranges using the symbolic information generated by the compiler. Then you can correlate what's happening on the target system with the source code. You can't yet click on a symbol in the performance analysis window and bring up the corresponding source code in, say, a debugger, but Tektronix is working on automating the process.

Statistical, real-time analysis

Performance analysis on a logic analyzer, such as Tektronix's new GPX, has statistical analysis available using the DAS but also includes real-time analysis. Using symbolic information downloaded from the compiler's output, you can set up 12 real-time ranges and monitor up to four CPUs simultaneously—or, with a single CPU, up to 48 ranges at a time. Using both statistical and real-time analysis and starting at a high level, it's possible to identify areas where there might be problems and then home in on them.

Microtek International (Hillsboro, OR) provides what it calls a Swat board and software for its PowerPack 32-bit emulator; the device even does state histogram measurements, telling what percentages of time the CPU is spending in which routines, all in real-time. According to software design engineer Bruce Ableiding, the trace buffer



The Tektronix digital analysis system (DAS) gives insight into how well code is running at a structural and architectural level. One of its capabilities is code profiling, in which the display is divided into address ranges which show a statistical sample of what percentage of the CPU's time is being spent in which parts of the program. This lets developers home in on performance bottlenecks that might not be apparent from the level of bytes and registers.

analysis, on the other hand, catches every event as it happens.

Statistical performance analysis is used, for example, to monitor a set of address ranges on a bus and to record the percentage of execution time spent in different ranges. Because the logic analyzer is reading bus activity, updating counters and managing the display, it won't see every event that takes place. But if the program is run long enough, the analysis tool will gather enough data to give an accurate picture of what portion of the CPU's time is used where. "If you have the time to let your program run continuously," says Berger, "statistical perfor-

CPU cycles than it actually does. One way around this is to run the program several times with different sampling rates.

In real-time performance analysis you don't want to miss a single event, because you may be looking for that one anomalous instance where the system's performance goes out of spec. In this case, the analysis tool has to keep up with the speed of the system.

"In real time, the event doesn't pass into the buffer," says Tektronix's Berger. "The analyzer is seeing the sample coming through, immediately processing it and is prepared for the next one right



So many packaging systems choices.

The backplane is the single most critical element of a packaging system. Augat offers the broadest selection of quality backplanes available, many in stock.

Our VME prototyping tools continue our long-standing tradition of assisting the development engineer; high performance extender boards speed up critical debugging process and wire-wrap panels provide high-quality prototyping panels for designs using in high-speed devices. VME powered chassis make it easy to test new designs.

We offer a line of high performance VME and VXI backplanes equipped with an innovative 96 pin DIN connector with internal switching. The "Auto Bus Grant"

connector provides jumpering between the bus grant daisy chain and the bus arbiter when a card is removed from the VMEbus or VXIbus.

Send us your specifications and we'll quote your requirements on stock, custom or semi-custom product.

We encourage you to get complete technical and delivery information on these or any other Augat components that you require. Call today for your free VME VXI Engineering Design Guide.

One solution.

AUGAT®
Quality
and Innovation
INTERCONNECTION PRODUCTS DIVISION

452 John Dietsch Boulevard
Attleboro Falls, MA 02763 USA
Tel: (508) 699-9800 FAX: (508) 699-6717

TECHNOLOGY FOCUS: PERFORMANCE ANALYSIS

of the emulator, running at the clock speed of the processor, can trigger a trace of an event, store that trace in a subbuffer and then trigger another subbuffer trace of another event. It can also measure the time intervals between such events. So, even when doing state histograms, the emulator is counting every event, and isn't just taking a statistical sample.

The problem with statistical analysis, according to Ableidinger, is that a program has to be sampled repetitively to measure the relative times you spend in one procedure versus another. But if something just occurs once, the statistical system might not pick it up. It's up to the developer, therefore, to determine how important absolute real-time measurements are in deciding what type of performance analysis to use. Whichever method you choose, at some point you'll probably want to move smoothly from the higher architectural level to a more detailed—perhaps even register-level—view of the situation.

Zooming in at the right level

Using performance analysis to locate bottlenecks is "a lot like [using] a zoom lens," says Applied's Maggard. "You need to slide up and down the scale from way up at the system-level view to way down at the nuts-and-bolts view." This is especially convenient when using performance analysis on an emulator, where you have control of the CPU and can easily refer to high-level source code in the debugger and can also examine registers if need be. The advantage of a logic analyzer is that it can monitor external signals in the target system that may point to where a problem lies.

Hewlett-Packard (Colorado Springs, CO) provides performance analysis options consisting of plug-in boards and software for both its HP-16500 logic analyzers and its HP-64700 series of emulators. Although the instruments have different uses and provide different advantages, John Marshall, manager for analyzer products at HP, says that "If you're using a logic analyzer, you wouldn't buy an emulator just to do performance analysis." Both performance analysis packages provide similar capabilities, including three major analysis modes: state overview, state histogram and time interval mode.

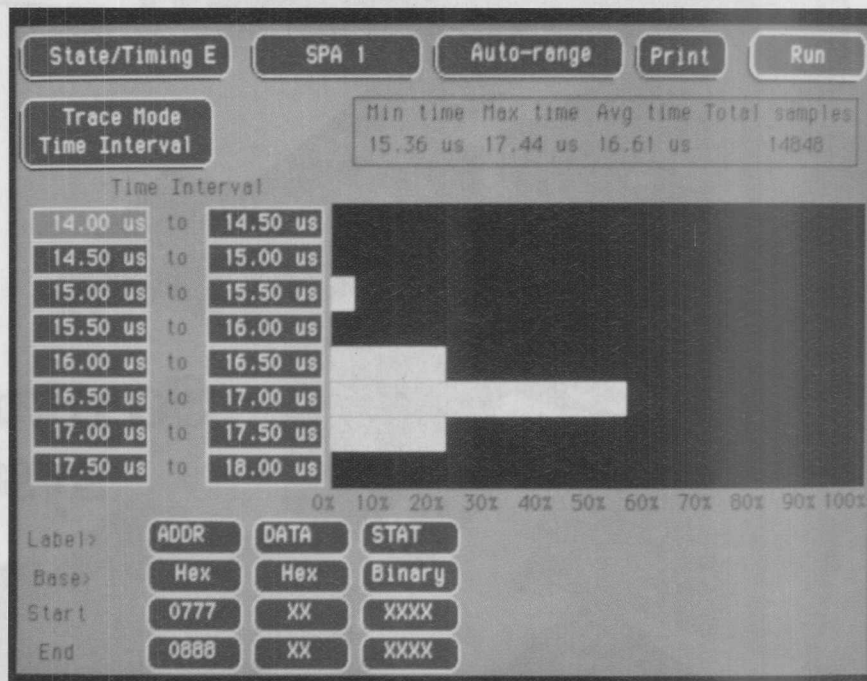
Normally, you start looking at system activity in state overview mode,

which simply monitors a contiguous range of addresses on a bus. This range of addresses is divided along the X axis of the display into 256 equal subranges, or buckets. The sampling produces histograms of how much time the code is spending in each subrange. After looking at the overview, you can select subranges that show unexpected levels of activity and take a closer look at them in the state histogram mode.

The address ranges specified for state histogram mode analysis can be non-contiguous and can be iden-

longer than specified, it may signal a problem. Both the logic analyzer and emulator versions of the performance analysis system can trigger on an overtime. The logic analyzer can then display a trace of events leading up to and following the event of interest—to see, for example, if some external I/O problem is to blame. The emulator can also trigger a trace and export a signal to an external logic analyzer, which can be monitoring signals on the target unavailable to the emulator.

While emulators can examine ad-



The time interval mode display of Hewlett-Packard's System Performance Analysis (SPA) option shows different execution times for a single routine. The routine is specified at the bottom by starting and ending addresses. The bars of the histogram show how many times the execution of that routine falls into selected time ranges.

tified either by address range or by name, using the symbolic information generated by the compiler. The state histogram reflects the ratio of the samples falling within the specified ranges to the total number of samples taken. A function that takes an inordinate amount of time or that interrupts the CPU more often than expected will stand out here, so that you can then look at the timing characteristics of a selected routine using the time interval mode.

Time interval mode shows the distribution of execution times for a single event. You can set up to eight time interval ranges, which don't have to be contiguous, and monitor the length of time it takes the routine to execute. If a routine takes

address ranges used by a real-time operating system (RTOS) just like they can look at other code, that doesn't mean they're operating-system-aware in the sense that you can easily look for specific RTOS tasks or service routines or can track dynamic memory usage. Integrated Systems (Santa Clara, CA) has teamed up with Hewlett-Packard to provide support for its pSOS+ operating system using HP's 64700 emulator series. This support includes performance analysis.

Software-only analysis

You can now specify the symbolic names of routines in your application that are to be followed with performance analysis. You can also

IN REAL-TIME GIVES YOU THE FREEDOM TO CHOOSE.

Whether you're looking for speed, flexibility, or compatibility, only VxWorks® provides this much freedom:

1. MULTI-PLATFORM. VxWorks gives you the freedom to choose from over 200 commercially available platforms, including all of today's most popular architectures.

TARGET CPUS

68K
SPARC
MIPS
i960
683XX


2. SCALABILITY. It gives you the freedom to develop applications of any size - small enough to control a fax machine, or large and complex enough to handle electrical power distribution for a major metropolitan area.




VxWORKS CONTROLS A
NEW LINE OF WORLD
FAMOUS LASER PRINTERS.

most widely ported real-time operating systems.

5. NETWORKING. VxWorks offers industry standard networking for real-time applications. To this day, it is the only real-time OS that gives you the freedom inherent in fully integrated networking.



A LEADING EUROPEAN
AUTOMAKER CHOSE IT FOR
PERFORMANCE AND RELIABILITY.



VxWORKS CONTROLS THE TRAFFIC
SIGNAL LIGHTS AT 12,000 OF NEW YORK'S
BUSIEST INTERSECTIONS.

3. PROVEN PERFORMANCE. VxWorks' high-performance kernel is fast, powerful, and easy-to-use. So you're free to concentrate on building your application and getting it to market quickly.

4. PORTABILITY. With VxWorks you have the freedom that comes from using one of the world's

6. TOOLS. VxWorks gives you

the freedom to choose from a suite of powerful development tools. And the freedom to use only what you need.

7. SUPPORT. All VxWorks customers have the freedom to call upon our worldwide team of specialists, who are there to help with everything from training to documentation to technical support.

8. ACCEPTANCE. It gives you the freedom and peace of mind in knowing that VxWorks has been chosen by major OEMs, such as Intel® and Motorola®, as one of their own real-time operating systems.


9. STANDARDS. VxWorks gives you the freedom to work with all the leading open system standards.

10. OUR MISSION. We are pledged to provide developers with software and service solutions by giving them the freedom to create real-time applications quickly, economically, and with minimal risk.



WIND
RIVER
SYSTEMS™

VxWORKS IS
FOUND IN
TELECOMMUNICATIONS AT
EVERY LEVEL
FROM PBX
SWITCHES
TO HANDSETS.



TO GAIN YOUR
FREEDOM CALL
1-800-677-1586

©1992 Wind River Systems, Inc. Wind River Systems and VxWorks are registered trademarks and the Wind River Systems logo is a trademark of Wind River Systems, Inc. UNIX is a registered trademark of UNIX System Labs, Inc. All other brand names and products are trademarks or registered trademarks of their respective companies.

CIRCLE NO. 45

TECHNOLOGY FOCUS: PERFORMANCE ANALYSIS

identify specific psos+ tasks and routines and can more easily examine how the application is interacting with the multitasking operating system.

While real-time performance analysis requires fast hardware, it's often possible to get very good statistical analysis using software tools alone, if those tools can be made sufficiently unintrusive. Software-only tools that can be used with simulators or CASE tools are also able to give you a look at the performance characteristics of your design before you've committed it to hardware. Instruction-set simulators such as the FreeForm/Simulator from Software Development Systems (Oak Brook, IL), for example, yield exact cycle counts that can be used to predict system performance.

Raw cycle counts in a simulation don't show you statistical variations that may arise, unless you're also simulating I/O and interrupt conditions. But they do give you an indication of, say, what clock speed you'll need to use for a given processor. With the FreeForm/Simulator, you enter parameters for processor and clock speeds and then run different combinations to target a desired price/performance point. You can also select critical routines by name and obtain cycle counts between two specified symbols or addresses.

"We consider FreeForm/Simulator to be less of a diagnostic tool and more of a front-end analysis tool for the design process," says Software Development's president, Jim Challenger. "It lets you get down to pieces of code and say, 'This is where I need to start optimizing my code.'"

Software-only performance analysis tools that monitor the target hardware need some way of gathering data other than by sensing electrical signals on the target's pins. This means some minimal intrusion into the target, but the monitors, or daemons, that are now being used by such companies as Wind River Systems (Alameda, CA) and Ready Systems (Sunnyvale, CA) are so un-

intrusive that they can perform very accurate statistical analysis. (The daemon, pronounced 'demon,' derives its name from a supernatural being in Greek mythology that stands intermediate between the gods and humans. In performance analysis, the daemon is a monitor that resides between the target system and the analyzing system.)

Wind River Systems, for example, offers such a tool called StethoScope as part of its VxWorks development environment. StethoScope runs on a Sun workstation under X Windows

or with mathematical analysis tools such as Matrix-X from Integrated Systems or Matlab from Mathworks (Natick, MA). ScopeProfile is, as the name implies, a profiler that shows how the real-time program is using CPU resources.

Since StethoScope operates under VxWorks, you can find out about kernel activity by name as well. "You can watch how many tasks are blocked on a semaphore or how many bytes are queued on a socket," says Wind River's product marketing manager, Catherine Jaeger. Since using Stetho-

Scope presumes you have running hardware, it's primarily used as an analysis tool for applications.

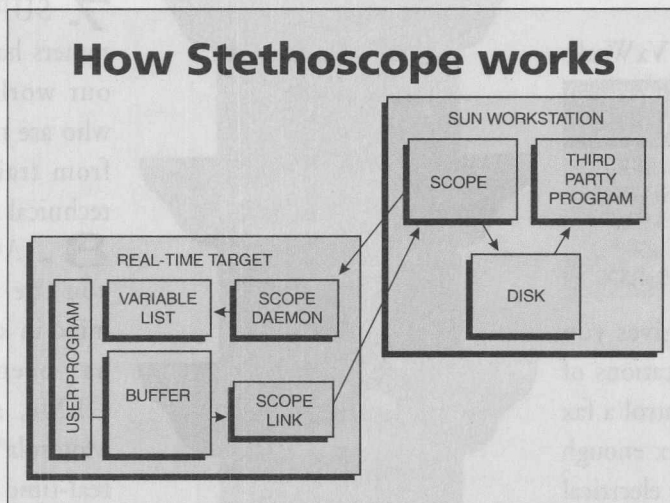
To dream...

What would be nice, of course, is an analysis tool that could characterize performance in the design stage and then monitor performance of the actual system once it was moved to target hardware. Ready Systems says it's responding to designers' wishes in this area. What designers are looking for, says Inbar Lasser-Raab, senior product marketing engineer for the company, "is a software analysis tool that can help during coding, not just during precoding."

Ready has a tool called VRTX Designer that can show performance behavior at the architectural level.

Another thing designers are looking for, says Lasser-Raab, is operating system awareness that isn't limited to a single kernel. They want to measure their interrupt service routines, because these are often the most critical parts of a real-time application. And they want to see how task switching occurs, so they can look for things like starvation and lock-out—situations in which a task is unable to get sufficient CPU time or is blocked from executing at all. "People really prefer software-based tools because they are easier to use," Lasser-Raab says. "You don't have to deal with all kinds of cables."

Since such a tool samples the running application and then displays the analysis results offline, it could theoretically be used to measure the




StethoScope, available with Wind River's VxWorks, runs on a Sun workstation and interacts remotely with the target via the real-time signal manager module, called ScopeDaemon. The target program places data in a buffer that is transferred to StethoScope via the ScopeLink cod, which is designed to have a minimal impact on target performance. A third-party program, such as Matlab from Mathworks, can be used to analyze and process numeric data.

and communicates with the target hardware system via a network. Two processes run on the target. A signal manager called ScopeDaemon processes commands and parameters and loads variables to monitor the program and perform calibrations. As the program executes on the target, it places data in a local buffer. A second target-resident process called ScopeLink transfers the data to StethoScope, which is running on the workstation where the data can be analyzed.


StethoScope's major analysis tools are ScopeProbe and ScopeProfile. ScopeProbe is a library of routines that can be loaded onto the target to collect time histories of variables in real time. Because the data is buffered, you can save the results of different runs for comparison. Data gathered in this way can be examined with StethoScope's own analysis tool,

Think Performance.

400,000 gates



Think Mitsubishi Gate Arrays.



Before you design your next gate array, or even your first, you've got to think about performance. Your very next thought ought to be: **Mitsubishi Gate Arrays.**

Mitsubishi's triple-layer metal, 0.8 μ m gate arrays offer 400,000 gates with over 60% utilization, and a typical loaded delay as fast as 215 picoseconds. We also give you up to 512 I/Os and pin counts as high as 576 in our exclusive μ Pitch TAB™ packaging.

We also offer design kits for the industry's most popular workstations, from logic synthesis, to simulation, to automatic test pattern generation (ATPG). So you can design on your own workstation or ours.

Mitsubishi offers both local design support and the global resources of a stable, well-capitalized company. As one of the world's top 10 semiconductor suppliers, we've been in the ASIC business over 15 years and we're continuing to invest in technologies for the next decade.

When you think gate arrays, think performance. Then think Mitsubishi. You'll be glad you did.

Phone (408) 730-5900 ext. 2106

VISION®
ACTION
REALITY

 **MITSUBISHI**
ELECTRONIC DEVICE GROUP

CIRCLE NO. 46



At LCC, a cellular systems engineering company, we created a field measurement tool based on a 68000 single-board computer running the C-Executive operating

system. One of the tool's requirements was that it save all measurement data to a hard disk. After our code was completed, we discovered that not every measurement was being saved. We knew we were inefficiently performing the tasks that had to be done, but we didn't know which ones were inefficient or how inefficiently they were running.

The creators of C-Executive, JMI Software Consultants (Springhouse, PA), offered a solution to our problem in the form of a profiler that uses the clock interrupt routine, which handles the operating system time slices, as a way of sampling which process is active at the time of the interrupt. The samples for each process can be accumulated and presented as a percentage of the total number of clock interrupts counted.

■ Optimization goals

Our primary optimization goal was to reduce the percentage of the dominating process. We also wanted to know the number of context switches accumulating from one process to another and from one process to the scheduler and back again.

A secondary optimization goal was to reduce context switching. A high-priority process that reads a data source a single byte at a time can produce extremely high numbers of context switches while not spending much time running the process. Even with optimized context switch times, the extra burden adds up quickly. During testing, we ran our firmware with particularly

difficult data loads for a constant test time and observed which processes were accumulating the most hits (represented by time used per process) and how many times each process was preempted by another process or by the scheduler. While the sampling process doesn't give accurate absolute values for time spent in each process, the relationships between time spent in processes are indicated. The inaccuracy is due to context switches that occur within consecutive system clock ticks and prevent sampling while the process is active.

Since C-Executive lets multiple processes run the same code (shared text), in-code markers normally can't resolve between processes. One advantage of process profiling is that text can be optimized after learning which process is the dominant user.

During our testing, process percentages showed varying rates of accumulation. This could have been caused by processes having high initialization loads but low steady-state loads. To see the current rate, we added a running rate of samples over the last minute for each process. This let the current rate of activity rather than the overall accumulated result be displayed.

■ The results are in

We expected the dominant process to be the interface to the measurement device. The results of our first profile were surprising; the display server was running at from 20 to 50 percent of the total hits. By reducing the number of separate writes to the display by a factor of 80:1, we managed to decrease the display server hits to under 10 percent. The extra time showed up as dead time in the scheduler, which increased from 20 percent to 50 percent. All of this was done in a low measurement data rate mode. After switching to a

high data rate mode, we then found the expected answer. The interface process was as high as 50 percent after the optimization we had performed on the display server.

Examining our code quickly explained why this was so. The interface was set up to read and write in an end-of-text mode, an efficient method that reduces I/O calls; an overriding parameter, however, limited the size of reads and writes to one byte. Rewriting the interface reduced I/O function calls by factors ranging from 20:1 to over 300:1, depending upon the data.

The profiler revealed the improvement. The interface process was now requiring only 10 percent.

Attention then centered on another process used to filter data and maintain an internal database, which the profiler showed was taking 29 percent of the time. But at this point our profiler reached its limit—we couldn't see which subroutines in this process were creating the bottleneck. JMI suggested that we place markers in each subroutine and read them in the clock interrupt service routine, extending the resolution of the profiler beyond the process limit. This then let us profile within our difficult process.

By reading queue sizes, we could tell if a process wasn't handling its data load. When we eventually found two processes that were instrumental in balancing the data load, we implemented a dynamic system of time allocation between them which was self-balancing.

Each attempt at optimization was immediately recognized in terms of its value. No effort was wasted optimizing where it wasn't needed, and failed attempts were immediately abandoned. After a week of optimization, we achieved our goal for saving all data, thanks to the quickness of the profiling process.

John Morfit, senior project engineer, LCC Incorporated, Arlington, VA

code running under a simulation as well as the code running on the target. It could characterize the system in the design phase—that is, help select the hardware execution environment—and then analyze the program on the target system. Ready Systems has created an environment that will support such a tool with its Spectra cross-development

system.

Spectra is a host-based tool interface and target manager that can monitor target activity with what the company calls the Xtrace Daemon, which runs on the target in somewhat the same way as Stetho-Scope does.

The difference is that Ready's Spectra environment also provides


for a virtual target that can run on the host workstation. Design tools such as debuggers, linkers, librarians, and test programs can work equally well with either the virtual or the real target. A performance analysis tool could potentially be used in the same environment.

Of course, software development doesn't have clear boundaries. If bot-

I THINK LOW POWER.

2.4 μ W/MHz/gate

Think Mitsubishi Gate Arrays.



Whether you're designing your next gate array, or your first, you've got to think about system power requirements. Your very next thought should be: **Mitsubishi Gate Arrays.**

Our 0.8 μ m arrays give you four speed/power options to control total chip power consumption. Four transistor sizes within each macro allow optimization for either high speed or low power. The result is power dissipation as low as 2.4 μ W/MHz/gate, at 5V. And, with Mitsubishi's 3V library, you can achieve even lower power dissipation. You can switch more nodes in

the array, control the power and still use lower-cost, plastic packaging.

Add to all of this 400,000 gates, 512 I/Os, and Mitsubishi's exclusive μ Pitch TAB™ packaging with pin counts as high as 576.

We also offer design kits for industry's most popular workstations, from logic synthesis, to simulation, to automatic test pattern generation (ATPG). So you can design on your own workstation or ours.

With both local design support and the global resources of a stable, well-capitalized company, Mitsubishi is one of the world's top 10 semiconductor suppliers. We've been in the ASIC business for over 15 years and we're continuing to invest in technologies for the next decade.

When you think gate arrays, think low power. Then think Mitsubishi. You'll be glad you did. Phone (408) 730-5900, ext. 2106.

 **MITSUBISHI**
ELECTRONIC DEVICE GROUP

VISION®
ACTION
REALITY

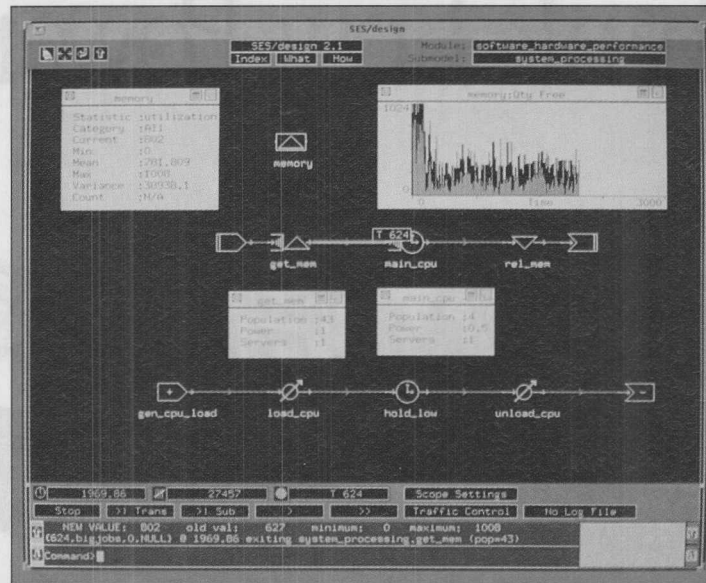
bottlenecks only show up when certain external conditions arise—such as when there are too many interrupts at once or some rare I/O condition occurs—then you've got to be sure you check for those conditions. That's where the problem bleeds over into the realm of software test automation, where scripts attempt to exercise code for all possible conditions that could cause a failure—perhaps something as simple as failing to meet a scheduling deadline.

It's important, then, in the early characterization of a system, to consider the workload, as well as how the software generally performs. "For example," says Michael Turner, vice-president of marketing for SES (Austin, TX), "you really need a tool that can tell you how much memory a process is going to take. If you don't use enough memory, you'll do disk swapping and that'll kill performance."

The memory solution

The solution for many people is to use as much memory as there's room on the board, but that can be less than cost-effective. "Workload is another real key," Turner points out. By characterizing the workload, you can get some idea of how much performance margin to build in once you know the system will work in a structural or an architectural sense. SES provides a tool called SES/Workbench that models the execution environment of a system in the design stage. Models are built up using four types of nodes that are represented graphically. Each node has a parameter list, so it can be characterized precisely in terms of such features as size and availability.

Resource management nodes create, allocate and release resources. Examples are processors, memory, communication links, and system buses. Transaction flow control nodes create, destroy and alter the flow of transactions through the system. Examples are loop, source,



This screen shows an SES/Workbench model of a system in which different jobs are competing for processor and memory resources. The upper right box shows memory use, while the icons represent the constructs that route the jobs. Parameters (list boxes) for each icon specify availability, size of queue, timing features, and other parameters that control the flow of jobs. With this model, you can increase memory size or processor speed or specify some combination of the two to improve performance, and you can see the results in the model before committing to hardware.

sink, and branch. Submodel management nodes let models be designed hierarchically. They can be exploded to reveal underlying nodes. Examples are reference, enter, return, and call. Miscellaneous nodes let you embed user-defined C code in the model.

SES/Workbench can be used in conjunction with CASE tools such as Software Through Pictures from Interactive Development Environments (San Francisco, CA). This makes it possible to use the software and execution environment models together so you can start modeling at a very high level. You can then take the model down to finer levels of detail. As code is developed, you can add parameters that exactly model the queues and delays of a given processor. Often, the significant bottlenecks are revealed at the higher levels of abstraction, so basic architectural decisions can be made before you've committed to hardware.

Modeling workload

You can also model the expected workload for the system. SES' Turner says there are several ways to do this. "On one extreme you can characterize statistical workloads. The other extreme is to do a trace on some real workloads from some com-

parable system and dump those into the Workbench model." Turner says that people tend to "grab the latest superstock 68040 at some astronomical clock rate, but when they get into the actual architecture they find they'd have a faster, more reliable system if they used three 68000s at a lower clock speed, because they find there's a lot of parallelism inherent in the design."

That this kind of discovery can come from performance analysis is reinforced by Tektronix's Berger. In the case of a system with two processors that were communicating with each other, his team was able to identify an algorithm that was burdening one processor while keeping the other from starting on its tasks. By changing the algorithm of the offending routine, he says, "We were able to reduce the

actual execution time by 64 percent." Such things are examples of the architectural and structural problems, not easily found with a debugger, that a step back from fine detail to a higher-level view can often reveal.

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

| | |
|---|------------|
| Applied Microsystems (206) 882-2000 | Circle 211 |
| Hewlett-Packard (719) 590-1900 | Circle 212 |
| Integrated Systems (408) 980-1500 | Circle 213 |
| Microtek International (503) 645-7333 | Circle 214 |
| Ready Systems (408) 736-2600 | Circle 215 |
| SES (512) 328-5544 | Circle 216 |
| Software Development Systems (708) 971-8170 | Circle 217 |
| Tektronix (503) 629-1925 | Circle 218 |
| Wind River Systems (510) 748-4100 | Circle 219 |

Creating Real Time

CREATIVITY TRANSLATES IDEAS INTO ACTION. IDEAS ARE NOT ENOUGH.

CREATIVITY IS CONTROL. IT CHANGES CHAOS INTO ORDER.

CREATIVITY IS VISUALIZING THE OUTCOME OF A DESIGN SOON ENOUGH AND CLEARLY ENOUGH TO IMPROVE THE PLAN BEFORE ACTION.

CREATIVITY IS THE EXTRA TEN PERCENT THAT MAKES ALL THE DIFFERENCE.

pSOSystem. Giving You the Freedom to Create.

The real-time operating system you choose for your next project can make or break it. You need a system that's bullet-proof and fast, integrated and complete. You want a system that gives you the freedom to focus on your application, to develop ideas into reality.

Give yourself that creative edge with pSOSystem, the benchmark real-time operating system. Scalable. Modular. One that's embedded and proven in thousands of products large and small, from telecom switches to video phones, from CT scanners to handheld terminals. And pSOSystem brings you the latest features you need:

- Open system networking, from NFS, RPC, to X11.
- Broad host support, from workstations, PC to VMS.
- Complete language support, from C, C++ to Ada.
- Integrated source and system-level debug tools.
- Transparent multiprocessing.
- Quick-start framework for custom or off-the-shelf hardware.

Discover creative freedom. Call us today at 800-543-pSOS, 408-980-0400 (fax), or at scg_sales@isi.com (e-mail).





©1992 Heurikon Corporation

Your success depends on having a firm grasp on today's performance choices. Technology, however, is only the beginning. Effectively integrating those choices demands performance from hardware and software alike. You should have the flexibility to select from standard or custom configurations.

Heurikon provides you with total solutions. Whether you use VME, Multibus or a proprietary platform we deliver you the real-time performance for your application. With Heurikon you have more than a grasp on the problem, you've got a handle on the solution!

With Heurikon
Integrated VME
Systems...

Get
A
Grip
On
Real-time

Solutions

HEURIKON[®]
A COMPUTER PRODUCTS COMPANY

Heurikon Corporation
8310 Excelsior Drive
Madison, WI 53717-1938
800 356 9602

Heurikon Europe
+44 (494) 88 32 32

Windows demands drive PC chip set features

PC chip set makers are focusing on Windows graphics issues. How well they handle acceleration, cache and power management will decide whether or not they prosper in the OEM market.

Don Tuite, Senior Editor

The big issues today for makers of PC chip sets and their OEM customers derive from Windows and its GUI. They boil down to graphics acceleration, cache and power management. These issues take on different weights, depending on whether the platform is a desktop, notebook, sub-notebook, or pen-based consumer appliance.

"When will I see what I'm going to get?" is a common cry of frustration from users of Windows applications. As a result, every chip set maker reports strong designer demand for faster graphics manipulation under Windows.

For the user, faster graphics involves a subjective judgment. In one simple demonstration on a test machine in the lab, an engineer grabs a face card in the Windows Solitaire application, drags it to a distant corner of the screen and lets it snap back into position. How readily the card follows the arrow cursor and how rapidly it snaps back give some idea of how quickly the system processes graphics.

A more realistic (but still subjective) test is to scroll around a word processing document that contains a lot of scanned-in, bit-mapped graphics. If you compare one of last year's 20-MHz 386SX notebooks with an accelerated 486DX desktop of today, you'll see that the document scrolls much more quickly and smoothly.

For an objective measure of Windows graphics performance, most chip set makers have accepted *PC Magazine's* Winmark, which is a geometric mean of eleven Winbench graphics benchmarks. For reference, the year-old notebook mentioned above, which snaps the jack of spades back into position with such alacrity, but which scrolls Microsoft Word files with embedded graphics with agonizing slowness, scores just over one million Winmarks. Makers of graphics accelerators, such as Weitek (Sunnyvale, CA), anticipate 20 to 30 million Winmarks for their Windows graphics accelerators. As an alternative to the *PC Magazine* Winbench suite, Microsoft (Redmond, WA) is developing its own MS-Test benchmark tools for Winstone benchmarks.

How an OEM goes about implementing accelerated graphics depends on



At Chips and Technologies, engineers Jill Seaman and Sunder Velamuri work with a next-generation flat panel. New, high-resolution color and monochrome displays, along with the Windows GUI, have created a demand for new laptop chip sets and VESA-standard local bus-based accelerator cards.

TECHNOLOGY FOCUS: PC CHIP SETS

whether the platform is a desktop or a notebook. In notebooks, which have no room for expansion slots, the graphics accelerator is part of the chip set and is necessarily installed on the motherboard. On desktops, however, OEMs have to decide whether their customers will prefer a motherboard-mounted accelerator or a graphics accelerator card in a slot. The first approach leaves all the card slots open, but locks in the OEM's choice of accelerators. The latter option lets users choose display modes and make upgrades, but ties up a slot.

In the desktop arena, chip set makers are seeing strong demand for card-based accelerators. Some of this stems from the approval this spring of the Video Electronic Standards Association (VESA) Local (CPU) Bus standard. Using the Local Bus for graphics removes the bottleneck between the CPU and the display processor when using the 16-bit, 8-MHz ISA bus.

Until VESA members agreed on a standard, chip set makers were either warning OEMs away from Local Bus or were touting their proprietary solutions. Today, that situation has changed. In fact, William Wong, product marketing manager for Oak Technology (Sunnyvale, CA), says, "I expect VESA is a great chance for cards to come back again. At least in desktops, we will see less of the motherboard-mounted graphics controller."

Inside an accelerator

According to Scott Cutler, vice-president of software technology for Chips and Technologies (C&T—San Jose, CA), many Windows accelerators consist of an SVGA controller with an interface to Local Bus, a built-in bit block transfer (Bitblt) hardware assist and single-port DRAMs for memory. He says that the single-port DRAMs establish the critical path for speed in these parts.

C&T's Wingine, Cutler says, is typical of higher-end accelerators based on dual-port VRAMs. When Wingine is running Windows, the host CPU updates video data through the

VRAM parallel port, while the Wingine accelerator processes data from the VRAM serial port to the display. The advantage of using this approach is that an increase in processor power manifests itself as an increase in graphics performance. (Parenthetically, when the processor isn't running Windows, Wingine

screen (frame buffer), screen-to-memory and, most useful in an accelerator, screen-to-screen. Screen-to-screen Bitblts are used for window moves, text scrolls and text line shifts.

After Bitblts, line draws are the most common function found in hardware accelerators, followed by

color-expanding pattern fills. Seldom offered functions include the abilities to handle arc, curve and circle primitives and to perform area fills.


Bob Payne, manager of graphics development at VLSI Technology (Tempe, AZ), says that roughly 80 percent of all application needs are covered by Bitblts, another 10 percent by line draws and about 10 percent by hardware assist features. In addition to these functions, Payne recommends that you get hardware cursor support, which places cursor data into the pixel data stream without touching the pixel data stored in display memory. The hardware uses horizontal and vertical position registers to locate the cursor within the displayed image.

The issue of VRAM versus DRAM goes beyond the number of ports affected. Chip makers agree that, compared to VRAMs, DRAMs are faster in terms of raw access

time, as well as being less expensive and available from multiple sources. Makers of high-end chip sets and VRAMs, however, such as Texas Instruments and Micron Technology, insist that the raw access time of DRAMs simply can't provide the required data bandwidth at high display frame rates, pixel densities and number of colors. Consequently, the consensus is that DRAM graphics memory is for low- and mid-range systems; VRAM is for the high end.

Bus issues

Mechanically, the VESA Local Bus header mounts on the motherboard just behind the standard ISA edge connector. A VESA display controller,



The most common benchmark among chip set vendors is the Winmark suite. These numbers, from Chips and Technologies, show over 20 million Winmarks of performance for a 486DX2-based system with a Wingine accelerator.

functions as an SVGA controller via the ISA bus.)

NCR (Colorado Springs, CO) goes a step further with its 77C22E+ VGA controller. The company adds a modest number of Bitblt hardware functions to a dual-port, Local Bus-based VRAM accelerator.

In general, graphics accelerators that include hardware support for Windows graphics contain functions such as clipping, Bitblts, line draws, curve draws, arcs, area fills, color expansion, and pattern fills. How many functions to support is another decision the OEM must make.

All graphics hardware accelerators provide Bitblts, of which there are three kinds: system memory-to-



© 1991 Brooktree Corporation.

Just What Your Customers Need, Another Outlet For Their Creativity.



What's in? Video Out. Putting video to a VCR and displaying video on a composite monitor are the newest capabilities every computer will need to compete in the Multimedia Age.

Now you're just a single chip away from adding Video Out to your very next computer design. Introducing Bt858, a monolithic digital device that packs in a board full of analog circuitry and puts out studio quality composite video.

Bt858 is a tweakless all-digital chip that bridges the video gap between RGB computers and composite or S-VHS outputs in the NTSC/PAL formats. It accepts multi-format digital inputs from 24, 16 or 15-bit RGB, 24 and 16-bit YCrCb and 8-bit VGA.

And because it has a programmable clock rate it adjusts for the 1:1 square pixels in computers and 4:3 rectangular pixels on TV without distortion.

Bt858 gives your system an image quality advantage, too. Studio quality output is a step

above tape decks and TV monitors so images always look "first generation."

You've read the book. Now see the picture. Call **1-800-VIDEO IC** and we'll send you "The Ins and Outs of Video Out," a revealing presentation of Bt858's capabilities.

That's all folks.

Brooktree Corporation, 9950 Barnes Canyon Road, San Diego, CA 92121, (619) 452-7580, FAX (619) 452-7294

Brooktree®

small package dual in-line package.

Intel's (Santa Clara, CA) SL version of the 386 (and, presumably, of the 486) integrates a great many functions on-chip, making the Local Bus inaccessible. The company's Peripheral Interface (PI) bus is a proprietary alternative to VESA. At this time only Intel and C&T have announced glueless PI interfaces to display controllers. But Intel has said that more PI interfaces will be announced this year.

A side effect of the growing popularity of VESA for graphics acceleration in desktops will be slower growth for EISA. All the chip set

sets to enter EISA desktops populated mostly with ISA boards, with just one EISA controller card for the network."

■ Simultaneous LCD/CRT displays

Another notebook/desktop issue is whether a motherboard-mounted display controller will support an LCD or a CRT. C&T has recently expanded its Vampire family of VGA controllers with the 65520 and 65530, which will support simultaneous LCD and CRT displays (as does Oak Technology's OakNote). Simultaneous displays let you control color projection displays from notebooks.

Providing simultaneous LCD and

essentially asynchronously, so the images on both displays are steady and flicker-free. Both types of chip sets require external buffers when OEMs use the simultaneous CRT/monochrome LCD capability; C&T's buffer is intended for VRAMs, Oak's and Cirrus Logic's are for DRAMs.

Chip set makers may also supply drivers for bigger cursors on the LCD display, for a portrait aspect ratio for pen-input appliances and for the ability to zoom into the main Windows screen and pan around it using mouse moves. This last capability helps you take advantage of high-resolution video modes on notebook screens.

There is a market for accelerator boards intended exclusively for flat-panel displays. C&T's Vampire chip set was designed into Yamaha's Display Master board, which targets the industrial systems integration market. One of the advantages of the C&T 65530 high-end, 1280 × 1024-pixel graphics controller is that it can be programmed for a variety of monochrome and color flat-panel displays.

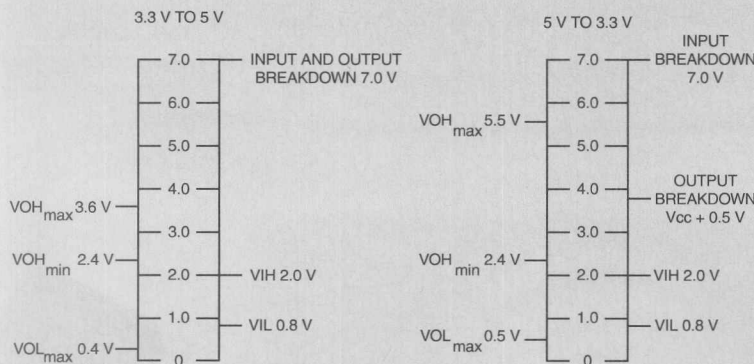
■ Cache and the marketplace

Graphics accelerators can't bear the full burden of making the Windows GUI run as fast as users demand. Much of that burden falls on the processor, which, at higher clock speeds, requires cache to avoid the DRAM access-time bottleneck.

It's a fact of life that, as long as Intel defines the PC processor architecture, chip set makers and OEMs must take into account Intel marketing strategies. The 486 introduced an internal 8-kbyte, four-way, set-associative cache and a four-stage write buffer. The line size is 16 bytes, with the processor bursting four double-words per burst. The burst mode is designed to provide page-mode DRAM accesses so that there are zero wait states on the last three double-words of the burst.

By including on-chip primary cache in the 486, Intel made it possible to increase the processor clock rate without speeding up the system clock, which is exactly what happens in the 486DX2 (for OEMs) and 486 OverDrive (for consumers) processors. As long as the cache hit rate is high, this translates into a substantial performance gain that an OEM or a designer can have simply by swapping chips.

JEDEC logic levels for 3.3- and 5-V devices

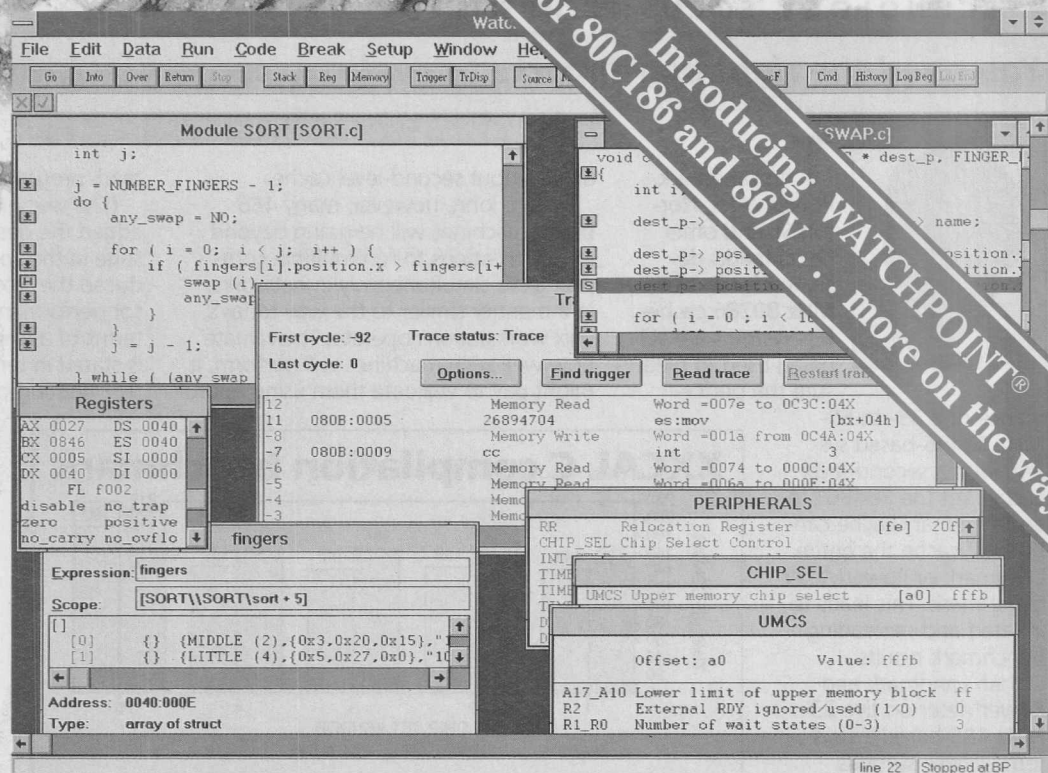


There's a need to interface between low-power 3.3-V hardware and high-speed 5-V hardware, but there are problems driving logic in either direction that are best handled by dual-voltage translators. A comparison of input and output standards for 3.3-V and 5-V parts helps show why some form of buffering between them is necessary. Using a 5-V device to drive a 3.3-V part directly will usually cause a latch-up failure as the ESD protection circuit in the lower-voltage device is overdriven. In the other direction, a 3.3-V device can nominally drive a 5-V TTL device. However, when the input stage isn't driven to the TTL rail, both the P- and N-channel transistor pairs turn on and conduct current approaching 1 mA. The resulting power loss defeats the purpose of using 3.3 V.

makers, whether or not they themselves are preparing any products for OEMs to use in EISA machines, see some future growth for EISA. They reason that the coming availability of OS-2, Windows/NT and some form of UNIX with a GUI will make EISA the bus of choice for servers.

Ron Mazza, vice-president of sales at Symphony Laboratories (Santa Clara, CA), which offers an EISA chip set, sees the cost of EISA coming down rapidly in the first quarter of 1993, as volume drives

CRT output is nontrivial since the two devices are scanned differently. While color LCDs scan essentially like CRTs, monochrome LCDs are usually dual-panel, with simultaneous scans of the top and bottom panels. In the past, there has been some discussion of using a FIFO buffer to delay the CRT data by half a frame and then interleaving that data, line by line, with the direct CRT data. This produces the required signals for the LCD's top and bottom panels, but at twice the CRT frame rate. Cirrus Logic (Milpi-




Windows Power Debugging ... for Real-Time Targets

WATCHPOINT® is a Microsoft® Windows™-based high-level language debugger for real-time embedded systems. In real time you can watch:

- ☐ Multiple source windows with user-controlled focus!
- ☐ Multiple data windows for variables, expressions, registers, memory, the stack, and trace data — with time-interval updating!
- ☐ Trace data, and you can examine and set memory, breakpoints, and triggers, *all while your program runs!*

Watchpoint gives you one-click navigation to routines, variables, and structures. The configuration tool bar puts the tools **you** need one

click away. And accelerator keys give you one-keystroke access to all common commands.



Breakpoint buttons appear on eligible lines. Set, delete, or switch between hardware and software breakpoints with one click.

Automate testing with unique full command capture and replay. Optional command sequence at break and trigger.

Watchpoint's Help system is the most advanced **anywhere**. Point to reduced screen images for instant help.

Complex triggers become simple with Watchpoint's Trigger Setup window. The informative Trace Display window differentiates reads, writes, source, trigger, etc., so you can quickly and easily find what you need.

Watchpoint takes *full* advantage of Microsoft Windows so you have easy-to-learn direct control over Sophia's powerful in-circuit emulators.

Call for more information.

Sophia
systems™
and Technology

1-800-824-9294

You need the right benchmarks to evaluate secondary cache



The Landmark, PowerMeter, Norton SI, and other benchmarks designed for 8- to 16-MHz 80286 cache-less systems are still being used to measure the perfor-

and without second-level cache.

Before long, however, many 486-based machines will be going beyond DOS applications to run multiple complex tasks simultaneously in networks, in a manner similar to the way today's UNIX workstations operate. To evaluate how well such machines will perform, it might pay to evaluate them using appro-

mark programs from the SPEC'89 suite.


(The SPEC group itself tacitly acknowledged the desirability of splitting the suite in the spring of 1992, when it introduced the SPEC'92 benchmarks. Processor performance is no longer given in terms of a single Specmark number, but is stated in terms of SPEC integer and SPEC floating-point performance.)

mance of a 33- to 50-MHz 80486-based system with a second-level cache. Yet the 80486's internal primary cache can by itself cache the entire Landmark or PowerMeter benchmark. This leads to inflated and misleading benchmark results.


I ran Landmark and PowerMeter on five 33-MHz 486-based ISA systems. These systems were built with chip sets from five different suppliers. The benchmarks were run without cache and with 128 kbytes or 256 kbytes of cache. Three of the systems had identical performance; two saw some marginal improvement from the addition of cache.

On the surface, this supports some chip set makers' skepticism about the utility of cache—at least in today's PC environment, where Windows 3.0 or 3.1 is used as a task-switcher with a GUI. However, there are benchmarks that show the effect of cache. I used the XXCAL suite created by XXCAL (Los Angeles, CA) to obtain the results in the top and middle sets of graphs to the right. XXCAL consists of five DOS applications run on Lotus 1-2-3, Dbase IV, Microsoft's C compiler, Rbase, and Word. The results from two of the benchmarks show marked differences among chip sets and between those chip sets with


XXCAL C compilation benchmark



XXCAL dBase IV benchmark



SPEC'89 integer benchmarks



appropriate benchmarks.

One possibility is to use a subset of the UNIX SPEC (Systems Performance Evaluation Cooperative) suite. I tested the same array of five 486-based systems using the four integer-intensive bench-

lower set of graphs. Again, there is considerable difference among the different chip sets, and the effect of second-level cache is striking.

The four benchmarks are Espresso, Eqntott, USP, and the GNU C Compiler. These are full applications that have run-times longer than one minute.

- **Espresso** is a routing tool for programmable logic arrays. It takes a Boolean function as input and produces a logically equivalent function that may have fewer terms. The benchmark measures the time it takes to run a set of seven input models.
- **Eqntott** tests sorting speed by translating a logical representation of a Boolean equation into a truth table.
- **LISP** measures the time USP takes to solve the 8-queen chess problem.
- **GNU C Compiler** measures the time it takes to convert 19 source files into optimized assembly language output. (This is Sun-3, not 486, assembler; the object of the benchmark is simply to measure compilation time.)

The results of using these benchmarks on the same 486-based systems are summarized in the

Ashok Bansal, president, AnR, Santa Clara, CA

crunch power. 30 MIPS and 5 MFLOPS change the definition of single board computer performance. Workstation features? This board is loaded.

If you need I/O, try 10Base-T, Ethernet AUI, SCSI, and VME. Parallel, serial, keyboard, mouse and audio ports get you connected.

Choose SunOS or VxWorks to tackle embedded and critical real time applications. Our MBus architecture is *the* high speed path for VME data and the only way to track Sun's hot new product moves.

Need more? Two SBus slots turn on a world of graphics and communications options. With 128 Kbyte cache and up to 128 Mbytes of DRAM on-board, no job is too big.

It's simple: When it comes to crunch, the 2SE eats the competition's lunch.

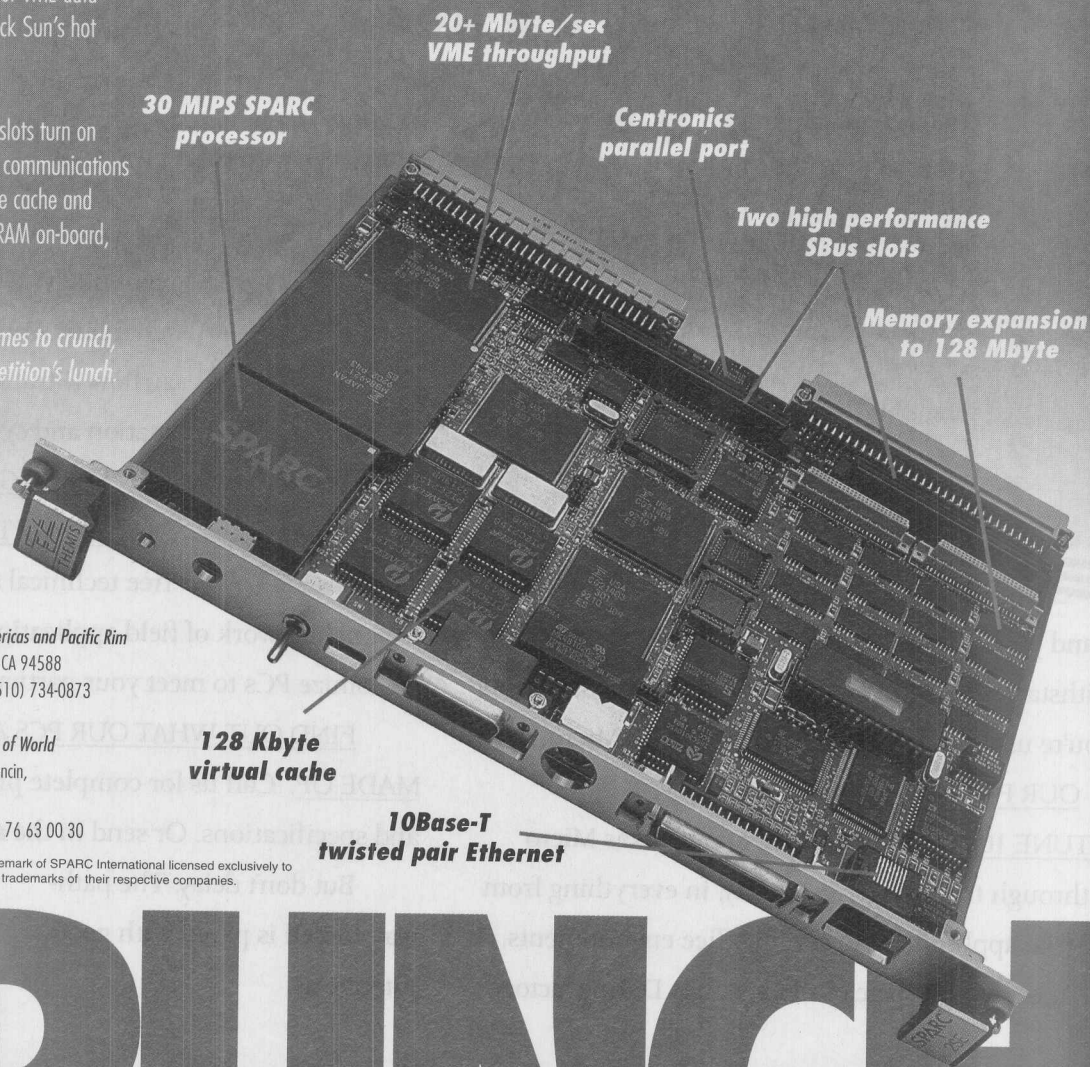


THEMIS COMPUTER • Americas and Pacific Rim
6681 Owens Drive, Pleasanton, CA 94588
Phone (510) 734-0870, Fax (510) 734-0873

THEMIS COMPUTER • Rest of World
1, Rue Des Essarts, Z.I. De Mayencin,
38610 Gieres, France
Phone 33 76 59 60 61, Fax 33 76 63 00 30

SPARC®2SE is a registered trademark of SPARC International licensed exclusively to Themis Computer. All others are trademarks of their respective companies.

30 MIPS 5 MFLOPS SUN COMPATIBLE



CRUNCH

IN A HARSH ENVIRONMENT, AN ORDINARY PC IS A DEAD PC. A lot of PCs do well to survive a desktop. But on the production floor or in the field, the dust, heat, vibration and traffic can finish an ordinary

tests, we shake, bake and beat them like there's no tomorrow. Because with an ordinary PC, there may be none.

OUR PCS ARE THE BEST BECAUSE THEY'RE BUILT FOR THE WORST. We design our 286, 386

and 486 systems using rugged design techniques that give them up to three times the life expectancy of other PCs.

Our passive backplane, for example, gives you instant access to plug-in CPU cards and components, reducing Mean Time To Repair to under 10 minutes.

We shock-mount the drives within our nickel-plated, all-steel chassis to withstand vibration. We implement VLSI and PAL technology to increase component reliability. And we use positive airflow filtration to reduce contamination and system heat.

STRONG SUPPORT IS OUR STRONG SUIT. We provide toll-free technical assistance and a


regional network of field application engineers. We also customize PCs to meet your particular specifications.

FIND OUT WHAT OUR PCS ARE REALLY MADE OF. Call us for complete product information and specifications. Or send in the attached card.

But don't delay. The pathway to hell is paved with good intentions.




If You Need A PC That Will Go Through Hell, Get One That's Been There.



PC—and your operation—in a hurry. Even if your PC has to withstand hell, your business doesn't have to. As long as you're using ruggedized PCs from Texas Micro.

OUR PCS TAKE A BEATING FROM THE FORTUNE 100. In fact, 70 of them put Texas Micro PCs through the wringer every day, in everything from industrial applications to severe office environments.

And we don't spare our PCs, either. During factory



Ruggedized Rackmount

- 14 option slots, 3 drive bays
- Nickel-plated, steel chassis
- 100,000 hr. MTBF* power supply
- Complete front panel accessibility
- 14" VGA Multisync monitor



Ruggedized Workstation

- 9 option slots, 2 drive bays
- NEMA-4 compliant
- 100,000 hr. MTBF* power supply
- CRT/CPU in one unit
- VGA monitor



Ruggedized Benchtop

- 10 option slots, 2 drive bays
- Nickel-plated, steel chassis
- 100,000 hr. MTBF* power supply
- Extra small footprint

Ask A Few Tough Questions: 1-800-627-8700.†

All trade names referenced are the service mark, trademark or registered trademark of the respective manufacturer. *Mean Time Between Failure. †European inquiries call 31-36-536-5595. Other international inquiries call 1-713-933-8050.

CIRCLE NO. 52

Although it's clear to OEMs that they can build products using the same motherboard and can differentiate them based on whether they have a 486DX or a DX2 processor, it's less clear that the OverDrive strategy will make sense to users. The reason is the confusing array of possible configurations. Intel has urged OEMs who offer platforms based on the 486SX to provide a socket for a 487SX math coprocessor. It turns out that this socket will also accept the 486SX OverDrive chip. Intel has also promised 486DX OverDrive chips, which you can install in existing machines that have OverDrive sockets. (Naturally, there would be no need for a math coprocessor socket with the DX.)

Regardless of processor, for both chip set makers and OEMs, consumer demand is pulling the cache bandwagon more than engineering and marketing departments are pushing it. Chip set makers and industry watchers note that common benchmarks show a modest improvement when primary cache is added to 386DX-based machines, or when secondary cache is added to 486-based machines. (But benchmarks of DOS applications may not tell enough of the story—see “You need the right benchmarks to evaluate secondary cache,” p 78.) Despite cache's questionable utility for the current run of applications, however, OEMs and their distributors report that users won't touch a high-end machine that doesn't have it.

While there are many controllers for external cache that follow the Intel model, some chip set makers do provide other options. Headland Technology (Fremont, CA), for example, has introduced the concept of virtual cache. Essentially, Headland engineers noted that the primary cache within the 486 handles reads very effectively. Product marketing manager Joe Nance says, “The primary cache and the burst mode of the 486 are very efficient. The hit rate is over 90 percent, and burst mode gets 75 percent of all reads with no wait states. It can get the other 25 percent with only one wait state.”

But cache write operations with the 486's simple write-through protocol are less efficient, says Nance. With write-through, there's a write to main memory every time cache is updated, and, he adds, “There's usually a minimum of one wait state on each write.” Headland says it can

improve memory write performance to zero wait states at a cost at the chip set level of under five dollars.

Nance explains, “We tried to provide a write buffer that empties faster than it's filled.” He says this is achievable because 486 write activity is less than 32 bits wide, and a large portion of the writes are to sequential locations. Since bytes are sequential, it's possible to accumulate them into 32-bit words before writing them to memory.

To make this work, the control logic has to handle cases where the

Regardless of processor, for both chip set makers and OEMs, consumer demand is pulling the cache bandwagon more than engineering and marketing departments are pushing it.



memory, to slow or stop the system clock when the processor is idle and to lower the operating voltage.

Western Digital's (Irvine, CA) WD7600 chip set for AMD's Am386SXLV processor is typical. There are separate timers for the LCD and its backlight and for controllers that power down the floppy and hard drives when they aren't being used. The WD7600 also shuts down serial and parallel ports when they're not in use. On the processor side, Western Digital chips use hardware control to slow or stop the CPU and to speed it up again, as well as to totally suspend the system except for slow CAS before RAS DRAM refresh.

C&T's Vampire graphics controllers, which can operate at either 5 or 3.3 V, have two power-down modes. In one, only the LCD is turned off; in the other, all operations are suspended and the controllers' outputs are tri-stated, while the VGA subsystem's contents are maintained using a 32-kHz clock. The Vampire chips can also switch off functional blocks when not in use, and can use reduced rates for panel shift clocks, vertical refresh and

processor must read from DRAM immediately after writing data to the buffer. This might happen, for example, when the processor pushes all its registers onto the stack and then reads code from a previously unaddressed routine.

The controller must also deal with the processor's preemptive reads, where processor reads have precedence over writing buffered data to DRAM. This includes the possibility of reading data that has just been stored in the buffer. Besides the Headlands controller, C&T's Winchips chip set offers the option of buffered writes to main memory in lieu of a conventional secondary cache.

Power struggles

In the past, power control was a goal only for notebooks. Recently, the Environmental Protection Agency has taken note of the number of PCs cranking away at ever increasing clock rates in countless offices around the country, all draining power and taxing building air conditioning systems to create further power drains. As a result, the EPA has offered incentives to manufacturers to reduce their products' power consumption and heat dissipation.

The first level of power management shuts down peripherals that aren't being used. These basic power management states are defined as: full-on, data entry, doze, and suspend.

- Full-on is equivalent to no power management.
- The Data Entry state assumes that an application is loaded and the user is entering data from the keyboard. The hard disk is not spinning and has been placed in low-power mode. Some CPUs, such as AMD's AM386 Battery Doubler, allow their clocks to stop between keystrokes.
- In the Doze state, the user hasn't hit a key in some time and, in addition to the power management conditions of the Data Entry state, the keyboard is turned off.
- In the Suspend state, all peripherals are powered down. Essentially, the only power being consumed is being used to keep DRAM refreshed.

Other ways to reduce the power consumption of a PC, beyond turning off peripherals when they aren't in use, are to use slow-refresh DRAMs for main memory and VRAMs for video

TECHNOLOGY FOCUS: PC CHIP SETS

video memory clocks.

Cirrus Logic's CL-GD62XX family provides similar functions and, in addition, can shut down operations entirely, saving operating states and memory contents to disk.

Mixing 3.3-V and 5-V operation

Today, the objective of using 3.3-V components is to extend battery life under full-on conditions to six hours and, under some form of power man-

agement, to eight hours or more. Because power consumption is proportional to the square of the supply voltage, cutting the latter from 5 V to 3.3 V reduces power use by about 56 percent.

Presently, few ISA peripherals and no disk drives or LCD drivers have been announced at 3.3 V. Moreover, there will always be some need to mate 3.3-V notebooks and palmtops with fast external devices that run

at 5 V. Consequently, OEMs must interface logic signals between 5-V and 3.3-V devices.

One problem is the question of whether CMOS devices run at CMOS or TTL levels. JEDEC specifies CMOS highs and lows relative to V_{cc} . Typically, on the input side, the CMOS logic level is 3.85 V, and on the output side, it swings rail to rail. This essentially precludes 3.3-V devices and CMOS parts that don't use TTL logic levels together. Furthermore, even though 3.3-V devices have the same nominal JEDEC TTL input and output specifications as 5-V parts, driving the lower-voltage chips from the higher can sometimes cause latch-up failures.


Bidirectional level translation can be handled by chips such as Quality Semiconductor's QuickSwitch, a by-10 array of MOS switches, Hitachi's HD151015 by-9 translator array with its own sleep mode, or IDT's 74FCT164245T, a 16-bit translator.

In the notebook world, there's been talk of running components at 5 V when the computer is plugged into an external supply and reducing the operating voltage of selected components to 3.3 V when the computer is running from its internal batteries. This lets the system run faster when it's using the external supply, and slower (and more economically) when it's on its own batteries.

Although it would be possible to create a system that would detect a change of power source, suspend operation, save the current operating state to disk, change the clock speed, and restore the operating state, this would be clumsy. An alternative method is to stop all I/O on the motherboard while a voltage change is occurring. Then, while the operating voltages are changed in small increments, the DRAM array is completely refreshed at each increment. This avoids a problem known as DRAM bump, in which large changes of V_{cc} cause data loss.

How low do you go?

Reducing the operating voltage raises several issues. One is just how far you lower it. Today, many chip makers take advantage of the ability of their standard CMOS devices to run with derated performance at 3.3 V, ± 5 percent. In addition to reducing power consumption, as the voltage swing decreases with the supply voltage, so does output slew rate. This helps reduce noise generated by the system, although

The advertisement features a large, detailed image of a Vigra audio board. The board is populated with various electronic components, including integrated circuits, capacitors, and connectors. A prominent feature is a vertical stack of connectors on the left side, including a large circular connector at the top and several smaller ones below. The board is shown at an angle, highlighting its components and the overall design. The background is dark, making the board stand out. The Vigra logo is visible in the top left corner of the board image. The text 'Real-Time DSP.' is written in a stylized font at the top left. The main headline 'Affordable real-time Digital Signal Processing' is in a large, bold font. Below it, the text describes the board's capabilities, including programmable coding and sampling rates, and mentions the 56001 DSP and multi-ported store-and-forward RAM. The bottom of the advertisement features the Vigra logo and contact information.

Real-Time DSP.

Affordable real-time Digital Signal Processing becomes a reality with Vigra's high-performance digital audio boards for VME and SBus. You'll find exclusive features such as programmable coding and sampling rates from high-fidelity DAT or CD performance down to telephone bandwidths (16-bit linear to 4-bit ADPCM).

Each board has a 56001 DSP for each audio channel, plus multi-ported store-and-forward RAM, on-board firmware in EPROM, and direct audio inputs and outputs. Vigra's unique DSP expansion bus provides an easy method for adding multiple channels and customized analog front ends.

For post-acquisition processing, use Vigra's **Audio Editing Toolkit** to make seamless splices and dazzling dubs — ideal for preparing interactive audio such as flight simulation, voice synthesis and communication response systems.

Vigra's board audio/DSP product line and custom capability assure you just the right match for your application. Call us today, and discover why Vigra is the leader in board-level audio.


VIGRA

4901 Morena Blvd., Bldg. 502, San Diego, CA 92117 (619) 483-1197 FAX:(619) 483-7531

CIRCLE NO. 53

CATS

COMPACTAT "CAT" INDUSTRIAL COMPUTERS



SMALL CAGE, BIG PERFORMANCE

Ziatech's large selection of CAT computers lets you build embedded systems to your exact requirements. The CAT's STD 32® architecture delivers ferocious performance in a variety of compact packages.

ON THE PROWL?

Call or FAX today for a free CAT data sheet

Phone 805-541-0488

FAX 805-541-5088

STD32®

©Copyright 1992 Ziatech Corporation. All rights reserved.

BIG CAT FEATURES

- 16- and 32-bit Industrial PC performance
- Industrial enclosures from 4 to 24 slots
- Multiprocessing, network and industrial I/O options
- Floppy, hard, solid state and PCMCIA storage
- DOS, Microsoft Windows®, OS/2 and UNIX compatibility
- STD 32 performance, STD-80 compatibility

ZIATECH
CORPORATION

Product names of other companies may be trademarks of those companies.

ply voltage requires ICs that were originally designed for 5 V to be derated is that, while internal device capacitances stay the same, the slower dv/dt reduces many critical performance characteristics. The challenge, then, is to build ICs that are really designed to operate at 3.0 V. This is the output of a pair of

computers. Making a true 3.0-V IC requires new IC designs and fab tweaks.

One example of a 2.7- to 3.6-V chip set that is available today is Headland Technology's 25-MHz PowerLite, used in AMD's AM386SXLV- and Cyrix's Cx486SLC-based systems. The two-chip set also contains 44 ac-

A final power-management issue is whether the system-monitoring portion of the process is carried out by the BIOS or within the hardware. Putting power management inside the chip saves development time for the OEM. Manufacturers are implementing power management either directly on the processor (Intel and AMD), or on the keyboard controller, usually the only other chip on the board with a microcontroller. Typical of keyboard controllers with integrated power management are Hitachi's H8/3332 and USAR Systems' UR5HCFJ8 Keycode for Fujitsu's FKB7211 notebook keyboard. Bucking the latter trend, Headland integrates power management in its 3.3-V HT25 single-chip core logic controller, rather than its HT35 peripheral controller. ■

Innovative DSP systems ...only from Pentek

How does Pentek do it?

By offering you a wide choice of baseboards and expansion modules that utilize the high-speed 32-bit Intel MIX bus...so data transfers and DSP performance are not compromised by traffic jams on your VMEbus or Multibus.

The right hardware...

- ◆ TMS320C40, C30, C25 and DSP32 engines
- ◆ A/D's and D/A's from 12 to 18 bits
- ◆ Sampling rates to 10 MHz
- ◆ Digital I/O and SCSI interfaces
- ◆ T1/CEPT telecom interfaces
- ◆ Precision clock generators

The right bus...

- ◆ VMEbus, Multibus I or II baseboards
- ◆ MIX bus expansion modules
- ◆ Bus adapters for SUN and PC-AT

The right software...

- ◆ Pentek SWIFT C development environment
- ◆ SPOX real-time OS and DSP libraries
- ◆ Comdisco SPW DSP code generator
- ◆ SUN UNIX and PC-AT Support
- ◆ ProNet Ethernet system
- ◆ DSP C compilers
- ◆ Drivers for MIX modules

Got a problem? We have the solution.

Call us today at (201) 767-7100

PENTEK

55 Walnut Street • Norwood, NJ • 07648
Tel: (201) 767-7100 Fax: (201) 767-3994

© 1992 Pentek, Inc.

System No. 1 in a Series

Pentek 4283

'C30 DSP
Processor
for VMEbus



Pentek 4247

Dual 'C30
Co-processor
MIX module




Pentek 4245

32-channel
16-bit A/D MIX
module (2 ea.)



System No. 1

64-channel,
98 MFLOP
signal analysis
system in 3.2"
of VMEbus
cage width



For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number.

Chips & Technologies (C&T)

(408) 434-0600 Circle 221

Cirrus Logic

(408) 945-8300 Circle 222

Headland Technology

(510) 623-7857 Circle 223

NCR

(719) 596-5795 Circle 224

Oak Technology

(408) 737-0888 Circle 225

Symphony Laboratories

(408) 986-1701 Circle 226

VLSI Technology

(602) 752-8574 Circle 227

Western Digital

(714) 932-5000 Circle 228

Weitek

(408) 738-8400 Circle 229

"FUZZY LOGIC" REPRINTS

Last year our March special report on **Fuzzy Logic** generated over **4,100 reader inquiries** and the **April 1992 Fuzzy Logic report** is already poised to generate an even higher number of inquiries! **Reprints of the April article are available now**—quantities of 100 or less are \$1.00 each, quantities over 100 are .75 each. To order, call Patti Kenney at (508) 392-2124.

**COMPUTER
DESIGN** Technology
and Design
Directions

CIRCLE NO. 55

FLUKE®



PHILIPS

Test results:



90% of those who try
a Philips Logic Analyzer from Fluke buy one.



100% get a free DMM*.

Our logic analyzers sell themselves. All we have to do is get one in your hands. To make sure you do, we're giving you a Fluke DMM*, whether you buy our analyzer or the competition's. (See attached card for complete details).

Only the Philips PM 3580 family of logic analyzers give you *true* dual state and timing on up to 96 channels - simultaneously. All accessible with one probe and one keystroke. Which means no more dual probing or reconfiguration between state and timing. Or no probes at all if you use our boundary-scan test option!

*The top-of-the-line Fluke 12 in our newest DMM family. It combines a smart set of troubleshooting features in a new design that's exceptionally fast and simple to operate — with one hand. It's yours after our 30 minute demo, no matter whose logic analyzer you purchase.

All our analyzers feature 50 MHz state and up to 200 MHz timing speeds. As well as integrated state and timing triggering for fast debug of complex hardware and software problems. Plus broad μ p support like Intel's i486; i386; 80286; 80186/88 families. The MCS-96, 8051, and i960 families. And the Motorola 68040 to 6800, 68HC11, 68332/1, 68302, 68340, 56001, AMD's AM 29030, and TI's 320Cxx family.

The PM 3580 family of logic analyzers is priced from \$4495 to \$11,450 - about half the cost of comparable analyzers. What's more you can have them up and running in only 30 minutes.

Find out why the PM 3580 family of logic analyzers were the only ones cited for

excellence and innovation by *Electronic Design*, *EDN*, *Embedded Systems*, *Electronic Products*, and *R&D* magazines. Take the Fluke Challenge. The odds are 100% you'll be totally impressed.

For literature, our video or a demonstration, call **1-800-44-FLUKE**.

John Fluke Mfg. Co., Inc., P.O. Box 9090, M/S 250C, Everett, WA 98206-9090. U.S. (206) 356-5400. Canada (416) 890-7600. Other countries: (206) 356-5500. ©1992. All rights reserved. Registered T.M. of Advanced Micro-Devices and Intel Corp. Ad No. 00178.

FAST ANSWERS

FLUKE®

CIRCLE NO. 56

Pushing the Edge of Capacity... 1.65 GB in a 5.25-inch Form Factor.

Edge-to-Edge Performance

The DK516C-16 uses Hitachi's advanced proprietary technology to deliver 1.65 GB of capacity and a fast 13.5 ms average access time.

Its SCSI interface provides a maximum data transfer rate of 5.0 Mbytes/sec (synchronous), with a 256 Kbyte data buffer and read look-ahead cache.

Or, if you have an ESDI application, look into Hitachi's 1.54 GB DK516-15 with a 14 ms average seek time and a 2.75 MB/sec data transfer rate.

Edge-to-Edge Quality

Choose the DK516 and you get a drive backed by the quality and reliability of Hitachi—a \$54 billion company. Unlike other drive manufacturers, we design, build, and test all key components in-house.

For more information about the DK516, or any Hitachi disk drive, call 1-800-HITACHI.

Hitachi America, Ltd.
Computer Division, MS500
Hitachi Plaza
2000 Sierra Point Parkway
Brisbane, CA 94005-1819



HITACHI®

Our Standards Set Standards

Authorized Distributors:

CONSAN 612-949-0053
(IA, IL, IN, KS, KY, MI, MN, MO, ND,
NE, OH, Pittsburgh, PA, SD, WI)

GENTRY ASSOCIATES
800-877-2225 (AL, D.C., FL,
GA, LA, MD, MS, NC, SC, TN, VA)

R SQUARED 800-777-3478
(AZ, CA, CO, NM, OR, UT, WA, WY)

SIGNAL 800-228-8781
(CT, MA, ME, NH, RI, VT)

SPECIALIZED SYSTEMS
TECHNOLOGY 800-688-8993
(AR, LA, OK, TX)

**HITACHI (CANADIAN),
LTD. 416-826-4100** (Canada)

CIRCLE NO. 57

IDA

INSTANT DATA ACCESS (IDA) DIAL (617) 494-8338 DOCUMENT NO. 1034

Please see us at COMDEX Booth #1416

Is neural computing the key to artificial intelligence?

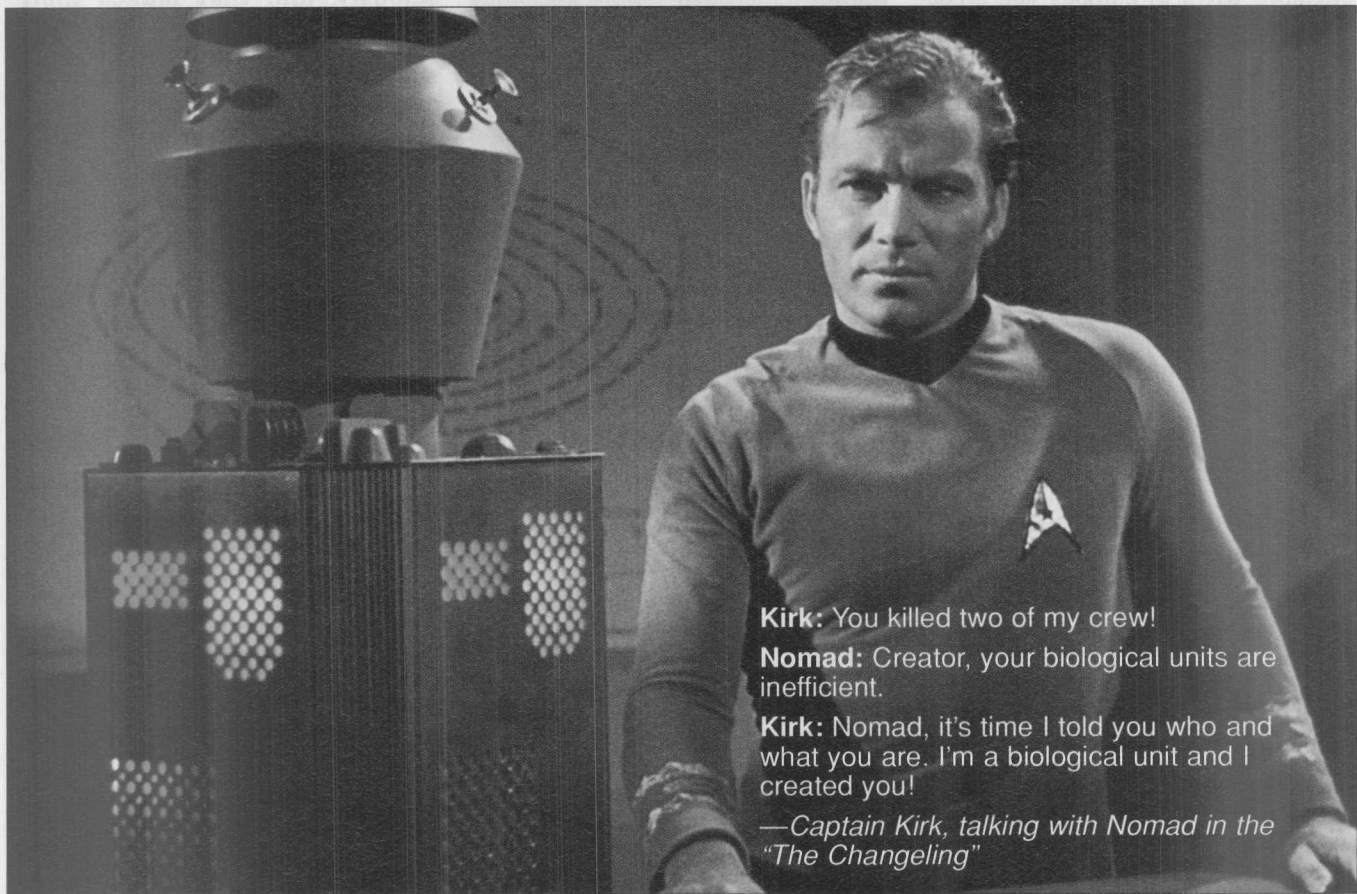


PHOTO COURTESY OF PARAMOUNT PICTURES

Kirk: You killed two of my crew!

Nomad: Creator, your biological units are inefficient.

Kirk: Nomad, it's time I told you who and what you are. I'm a biological unit and I created you!

—Captain Kirk, talking with Nomad in the "The Changeling"

Mike Donlin and Jeffrey Child

When people speculate about whether computers can learn to think, evil robots, such as Nomad running amok on the Starship *Enterprise*, often enter the conversation. In the *Star Trek* episode, Captain Kirk talks the sinister Nomad into blowing itself up, but the notion of a computer that can control human destiny makes many folks nervous. Computer professionals scoff at such silliness, even while acknowledging that advances in hardware and software have given computers the ability to emulate some human traits.

People have been fascinated with the idea of a machine that could think since the days of ENIAC, one of the first "electronic brains." ENIAC had 30,000 vacuum tubes and 50,000 relays, filled a large room and could rip through mathematical calculations at a blistering 13 operations a second.

Presently, computer technology resides somewhere between ENIAC and Nomad, but advances in artificial

intelligence, and particularly in neural networks, have caused a surge of interest in the thinking power of computers.

■ Nothing new about neural networks

The concept of neural networks has been around in some form since World War II, but it's only in the last six or seven years that working products have been developed that attempt to "learn" about and predict reality. In their infancy, neural networks and neural computing were the work of theorists who observed similarities in the way that computers and humans think. In both cases, a large amount of information is manipulated by breaking it into small particles—using gates in computers and neurons in humans. Gates handle data by fluctuating between an "on" and "off" state, and neurons do the same by firing (on) or not firing (off).

Scientists have postulated that human thought oc-

curs when two neurons fire simultaneously—and that their connection, called a synapse, is given more weight than would be the case if the neurons were connected but not firing. Because of these similarities between human and computer thought, researchers have begun to explore ways to embody the structures of human intelligence in machines.

In the case of connecting neurons, either through hardware or software, to emulate the brain, the task has been daunting. The human brain contains 100 billion neurons, each connected to 10,000 others by synapses. Building such a complex computer is a ridiculous idea, even with the staggering advances made in computer technology in the last twenty years.

Also, neural networks were dealt a blow in 1969 when Marvin Minsky and Seymour Pappert wrote a book called *Perceptrons*, which postulated that neural network research was a waste of time. Minsky, one of the founding fathers of the artificial intelligence movement, refused to believe that software could simulate the behavior of human neurons. Minsky's vision of artificial intelligence (AI) was far more comprehensive than just neural network technology, and he scoffed at those who wanted to reduce his broad theories to a set of equations that could solve only simple problems. Many experts blame this book for derailing neural network research and encouraging the expert-system theories favored by the authors.

Expert systems, in turn, have fallen out of favor in recent years, because they use a prohibitive amount of computing power to solve problems. Although there are some areas where encoding the skills of an expert and programming a



"It's really important to understand that neural computing has nothing to do with building brains."

—Casimir Klimasauskas, president of NeuralWare, Pittsburgh, PA



computer to carry them out seem feasible, for most complex tasks the intuition of an expert is simply too difficult to understand or too time-consuming to write out.

"To make an expert system work,

you have to know why the expert makes decisions," says Steve Bissett, senior vice-president at Synaptics (San Jose, CA), a neural network IC firm. "The problem is that most experts can't tell you all the rules that go into their decisions. They use past knowledge and intuition unconsciously, and that makes their knowledge difficult to codify. Even if you could write out enough knowledge to program into a computer, the amount of data would be so large that it would be prohibitively expensive to program and use."


Expert systems are by no means finished, but there's also been a renewed interest in neural computing in recent years, particularly in the fields of signal processing, forecasting and pattern recognition. Even though there are different ways to emulate the neural connection model of the brain, both in hardware and in software, all neural networks share certain common characteristics: they use artificial neurons that are connected to at least one other neuron, and they create their own representations of reality based on some form of learning model.

■ Learning by example

Fundamentally, all neural networks learn by association. For example, a neural network can learn to identify an apple by associating the inputs

"round," "red" and "fruit" with the output "apple." The neurons in a neural network are usually organized in three layers: input, hidden and output. Sometimes more than one hidden layer is used for complex analysis.

There are many ways that neural networks can learn, but the most common way is through example and repetition, also called back-propagation. Each time an input is given to the network ("round," "red"



In this simple neural network, a layer of inputs lists the various characteristics of fruit. When the network gets a stimulus from one or more of these inputs, it responds with an output. If the network guesses incorrectly, it adjusts the correlation of the internal connections or synaptic weights until it gets a correct answer.

or "fruit" from our example), it gives an answer. Naturally, when a network is new, the guess will probably be wrong. But over time, as the network gathers more and more data, it will begin to zero in on the right answer. When it's fully trained, it can deliver an answer that's more or less accurate, depending on the complexity of the task.

Every time the network guesses wrong, it adjusts its internal connections until it gets the right answer. These adjustments are made via synaptic weights, which give relative importance to data as it's applied to a task at hand. These weights can be implemented in hardware or software, and it's this ability to gauge the importance of

data that separates neural computing from a purely digital computational process. Although the synapses and weights can be made up of analog circuitry, digital components or software, the weighting procedure makes neural computing appear to have an analog nature—a characteristic that's especially important in performing tasks such as

Neural computing: What it is and what it isn't

Digital computers ushered in the information age, and as they have become smaller, more personalized and increasingly powerful, they've touched more and more aspects of our lives. But as brains for robots, computers still leave a lot to be desired. Recognizing natural objects in the real world, for example, is well beyond the grasp of modern technology. Are we just waiting for more computing power in a smaller package, or are we lacking a fundamental ingredient?

Computers perform logical computation. They operate on precise input information with a programmed sequence of instructions and produce a precise output. Computers are much better than humans at operations such as long division, yet, when it comes to pattern recognition, even insects process information better than the most powerful computer.

Scientists have been studying the brain for decades, trying to understand how biological computation works. Over the last ten years, a number of significant advances have been made and the embryonic field of neural networks has been born.

An artificial neural network is a computational structure similar to its biological counterpart, yet much simpler, even when compared to a very small portion of the brain. Nonetheless, artificial neural networks, simulated with digital computers, have already produced excellent results when applied to some real problems, such as predicting the outcome of horse races or playing backgammon. In many cases these results have been better than the best of the traditional logical or rule-based approaches.

Artificial neural networks perform what might be called intuitive computation. Rather than being programmed with a set of rules, they learn by example; they self-organize. A programmer


isn't required to figure out all the rules of the problem—a task that can be extremely difficult—and then write a software program embodying the rules. Rather, neural networks discover the rules for themselves through the process of training. A neural network contains a number of weighted parameters that are modified according to a learning rule, depending on the network's response to training examples.

Much of the research into neural networks has been performed using models built in software, simulating a neural network with a digital computer. While invaluable for research, this approach has limitations in any practical application that requires quick results, because the computer calculates the effect of each interconnection (or synapse) one at a time, whereas nature does it all simultaneously.

■ Biology, analog or digital?

Nature performs its computational functions using billions of neurons, each having thousands of synapses. To fully exploit the massive parallelism of the neural network structure, it's necessary to develop electrical circuits that approach nature's circuits in density, speed and power dissipation. Unfortunately, digital circuitry falls far short of the mark.

Analog circuits, on the other hand, can perform complex computations by exploiting the physics of semiconductors, which often mirrors the equations governing the behavior of biological neurons. In appropriate structures, this can lead to advantages of 100-to-1 in



function density and 10,000-to-1 in power consumption per function, compared with digital circuits.

The set of weights that embodies the learning of the network is stored in floating-gate structures; each analog weight value corresponds to an equivalent amount of electrical charge permanently stored in the floating gate. Synaptics calls this technology adaptive analog VLSI.

Despite their advantages, analog circuits can't match the density and energetic efficiency of nature, but the performance of analog circuits offers such a dramatic leap beyond digital circuits—and the precision of the required computation is so low—that they open up a vast new realm of everyday applications.

Synaptics solves pattern recognition and other problems by combining intuitive computation with logical computation. These two complementary forms of information processing are both essential in solving problems associated with autonomous-intelligent machines, that is, machines capable of naturally interacting with the real world.

Dr. Frederico Faggin, co-founder, president and CEO, Synaptics, San Jose, CA

Neural net chip speeds check reading and verification

Have you ever stood in line at a bank, the supermarket or a gift shop, waiting to cash a check and wondering why it takes so long? Chances are, it has something to do with waiting for a check to clear. Or maybe it's a delay caused by the cashier's inability to read a check.

To address this problem, Verifone (Redwood City, CA), a maker of transaction automation systems, asked itself how you outfit a check reader with enough computing power to handle any check, even those that are difficult to read. Checks in this category include those that are folded, wrinkled or improperly printed. Worse, ink density often varies from one set of characters to another. These variations can cause a great deal of trouble for check readers. Such machines either can't read the checks at all, or they read them inaccurately. Verifone wanted to make a check reader that could read 100 percent of the checks it sees.

Enter Synaptics, a company that was developing an application-specific neural network chip. To make the check reader project feasible, designers at Synaptics reasoned that a simple neural network application would be best suited for its chip. Because there are a limited number of characters on a check, 0 through 9 plus four special symbols, the Verifone application seemed an ideal candidate for a neural network.

Working with Verifone, Synaptics developed an analog neural network chip, the I-1000, specifically for reading checks. The design is historic in that it's the first commercial application using a neural network chip. In fact, the chip is so application-specific that it has a built-in lens through which it reads the images. The lens forms an infrared image on a "retina." The chip then tries to decode the character from this image. Making a number of computations, the chip comes up with an answer. It may see a "9," for example, assigning an 85 or a 95 percent probability to its answer, depending on how good the image is. A perfectly printed character with perfect ink density could receive a 99.5 percent probability. To allow for checks that are printed badly or are wrinkled or folded, Verifone trained the neural network to pass anything with a probability of over 80 percent.



You could find neural computing technology as close as your local gift shop. The Onyx check reader made by Verifone uses an analog-based neural network chip designed by Synaptics. With this chip, the unit can learn what a good check looks like, enabling it to handle a wider variety of checks—even checks that are folded, wrinkled or badly printed. (Thanks to the Forget-Me-Not gift shop in Auburn, CA.)

In an independent laboratory test, the neural-based check reader, dubbed the Onyx, was accurate 99.6 percent of the time. More important, it was capable of reading every check going through it.

Software and hardware support

While Synaptics was primarily responsible for the chip design, Verifone developed software and hardware to support the neural net chip. A 68HC11, Motorola's 8-bit microcontroller, controls all the processes in the machine and interfaces between the I-1000 and the software. While the I-1000 does the decoding, the Onyx also has neural network software. The job of this software is to make sense of what the chip tells it.

If the chip tells it, for example, that there's an 80-percent probability it's reading an "8," the software tries to determine whether the result makes sense. It captures frames to monitor timing as the check goes through the reader. Each frame runs for a different timing interval. The software determines if the acceleration of the frame makes

sense; it compares its conclusion to the result produced by the hardware. The final result is based on this comparison.

For memory, the Onyx has 28 kbytes of SRAM. Battery-backed SRAM is used instead of ROM to permit easy updates. If new types of checks are printed, Verifone can simply update (by retraining) the software in the company's lab. The software can then be downloaded to check readers over the phone lines.

Synaptics' neural net chip does more than just read the numbers, however. Using the magnetic properties of the ink used to print the bank number, the neural net can also determine whether or not the check is counterfeit. The ink used has a high iron-oxide content, and since most counterfeit checks are produced using a color copier, they wouldn't show any iron-oxide content. The infrared image indicates iron oxide in the ink based on the frequency of the emission received. If there isn't any iron oxide, no image is recorded, and the check is rejected.

pattern recognition or financial forecasting.

"There's a key difference between neural computing and digital computing," Synaptics' Bissett points out. "Traditional digital computing is like the left-brain or logical thinking that we do. The computer receives a set of rules or programs, then takes input and produces output based on those rules. By definition, it's only as good as the rules that guide it. Neural computing is more like right-brain thinking, which is intuitive. If you wanted a computer to read handwriting, you could try to write rules that would make it recognize an 'S,' for example, and that would be a very linear, but increasingly complex, way to solve the problem. That's not the way we do it. Our circuits are connected in parallel, and compare a shape to previous knowledge to try and categorize it. An important distinction, then, is to try and teach by example rather than programming by rules."

The correlation of neural networks to the way our brains work is what makes them suited to applications that need experiential learning, but is neural computing really thinking? In a word, no. Neural networks are patterned after the architecture of the brain, but in reality their ability to think is far more primitive than that of a common housefly. As a matter of fact, some experts scoff at the notion that neural networks are related to human thought at all, other than in a purely analogous way.

"It's really important to understand that neural computing has nothing to do with building brains," says Casimir Klimasauskas, president of NeuralWare (Pittsburgh, PA). "Neural networks are a collection of mathematical techniques that let you fit formulas to data, curves to data, and group types of data together. Neural networks could have been invented by statisticians, physicists or mathematicians, but the people who invented them were cognitive psychologists and neurobiolo-

gists, and so we ended up with the term neural networks. They have nothing to do with brains. I've found that if you try to explain neural networks from a human-thought perspective, people keep trying to fit them into a brain model, and it only confuses them."

■ Enter fuzzy logic


In spite of such caveats, most people will probably continue to associate neural networks with human thought, particularly because much

systems are converted into a neural network, which learns about the task and refines its knowledge. Once the neural network has achieved an acceptable degree of accuracy, it's translated back into a fuzzy system so its operation can be analyzed. According to Fujitsu researchers, this model reveals the hidden variables that develop in a neural network. By keeping the fuzzy rule sets and the neural networks as separate systems, the Fujitsu scientists believe they can learn more about how each system works.

But not everyone involved in combined fuzzy/neural computing wants to keep the disciplines separate. There's considerable activity in the AI community aimed at combining the neural network's ability to create relationships with fuzzy logic's capability to produce input and output information that spans a range of behavior. One such approach builds fuzzy operations into the learning techniques used by a fuzzy/neural controller. The resulting neural network learns to emulate a fuzzy controller, but with rules that can be altered by neural learning techniques.

In spite of the promise that the combination of fuzzy logic and neural networks holds, most of the work in this area is at the theoretical stage, although some practical applications have been demonstrated in medical imaging and flight simulation. To many theorists, however, the marriage of these two "smart" technologies seems inevitable.

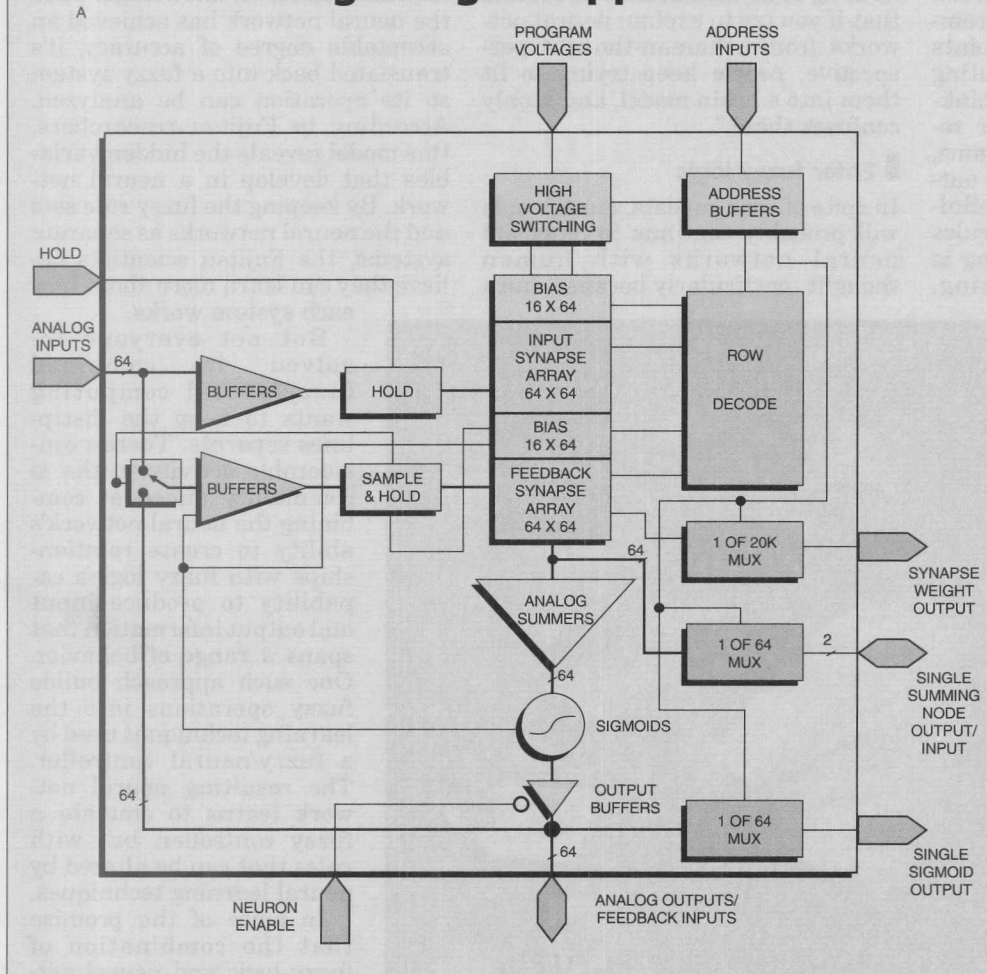
Theories aside, most of the tangible products using either fuzzy logic or neural networks have come from research that treats each of these disciplines as separate entities. In the field of neural networks, this means three categories of products: neural network ICs, whose architectures are specifically designed for neural computing; neural network software, which uses standard microprocessors and digital signal processors to emulate neural network behavior; and neural network systems, which are turnkey neural



"Neural networks have a lot of potential in optical character recognition," says Dan Hammerstrom (background), founder and chief technical officer of Adaptive Solutions, "especially where you have mixed fonts, uneven spacing and handwritten letters. In our OCR system, we point the camera at a page and isolate the characters from surrounding spaces and graphics. Our system then makes the characters uniform and passes them through a classification phase, where the letters are differentiated from one another and turned into ASCII code that the computer can use."

of the learning process in a neural network takes place in hidden layers or neurons, a processing paradigm similar to human thought. Some researchers are trying to unravel these hidden functions by using fuzzy logic techniques to better understand, or even work in conjunction with, neural networks. Fujitsu (Kawasaki, Japan) is working on a system, for example, that creates a fuzzy rule set based on "if-then" questionnaires that have been filled out by experts. These fuzzy

Analog & digital approaches

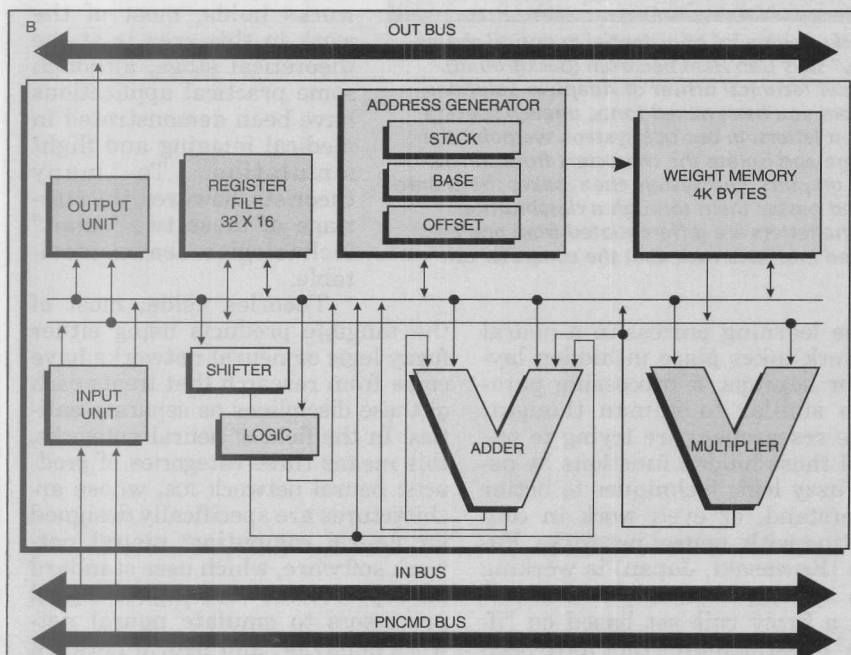


The 64 analog inputs to the ETANN device from Intel (top left) are capacitively stored for a limited time by taking the hold pin high. The ETANN's 10,240 synapses store weights as analog transconductance values, each producing an analog output current from an analog input voltage and a stored weight voltage. Currents generated by each of 160 synapses along a dendrite, or \pm column pair, are summed to form the net input to a neuron body. The dendrite's sum of currents is then converted to a voltage and passed through a sigmoid function with voltage-controlled gain. The CNAPS-1064 (bottom left) from Adaptive Solutions is an array of 64 processing nodes (PNs), each with 4 kbytes of on-chip SRAM. Each PN resembles a DSP and connects to three global buses: IN Bus, the data input bus; PNCMD, the command bus that dictates the operations of the PN each clock cycle; and the OUT Bus, an output bus.

network computers.

The first commercially available neural network-specific chip was the 80170NX electrically trainable analog neural network (ETANN) device from Intel (Santa Clara, CA). Introduced in 1989, the chip is a 64-neuron, 10,240-synapse IC with inputs organized as two groups, external and recurrent (or feedback); each input contains 80×64 synapse arrays.

Intel has also released a development system so that you can simulate, train and operate a high-speed neural network. Dubbed the Intel Neural Network Training System (iNNTS), the package provides two 80170NX devices, two learning simulation software programs, diagnostic software, a programmer interface, an adapter that can run on PC/AT-compatible computers, programming specifications, and full documentation. The iNNTS contains two learning simulation software




HOW MANY AMPS DOES IT TAKE TO POWER YOUR PRODUCT DEVELOPMENT?

One of the toughest jobs you face is analyzing and proving your initial design and performing board placement and routing. We'll take you from EDA to the MAX.

AMP DESIGN SERVICES:

methodologies, look into technologies and design pace with fast-changing envelopes, and generally keeping rates, coping with shrinking to market, dealing with higher clock If you're concerned with reducing time



FROM IDEA TO MARKET.

production experience in everything from metal and plastic to heavy machinery, precision optics to SMT boards to complete plug-and-play systems.

ment, and packaging design. And unparalleled tion and validation, systems develop interconnect design, analysis, simulation and connectors. They include Our capabilities today go far beyond connectors. They include trip to market easier and faster. companies. AMP can help make your schedules, can tax even the largest especially with today's tight development The resources you need for that journey

ONLY ONE.

S238B


AMP

S238/opc

FROM IDEA TO MARKET.

The resources you need for that journey, especially with today's tight development schedules, can tax even the largest companies. AMP can help make your trip to market easier, and faster.


Our capabilities today go far beyond connectors. They include interconnect design, analysis, simulation and validation, systems development, and packaging design. And unequalled production experience in everything from 'metal-and-plastic' to heavy machinery, precision optics to SMT boards to complete plug-and-play systems.



If you're concerned with reducing time to market, dealing with higher clock rates, coping with shrinking envelopes, and generally keeping pace with fast-changing technologies and design methodologies, look into what AMP has to offer, today.

AMP DESIGN SERVICES: EDA TO THE MAX.

One of the toughest jobs you face is analyzing and proving your initial design, and performing board placement and routing. We'll take you from




schematic or logic diagram all the way to a valid software first article.

Our comprehensive EDA toolset includes current, state-of-the-art software programs, and our own interconnect and device models—developed specifically for critical aspects of current design such as net simulation, timing verification,



Placement optimization uses thermal analysis, along with electrical, power, and manufacturing considerations.

(Courtesy of BETAsoft, Dynamic Soft Analysis Inc., Pittsburgh, PA.)



Characterized device and connector models are used in waveform analysis, timing verification, and board design.


analog transient noise effects, system impedance, and transmission line layout rules.

Talk to AMP early in your design cycle, and we'll put you on board faster—with flexible service offerings that range from help with specific problems to complete project management. Tap us for the years of experience we've put into engineering under critical electrical, mechanical, power, thermal,

and manufacturing constraints. And shave a lot of time-to-market and development cost off your project.

AMP FABRICATION SERVICES: MAKING IT REAL.


Major elements of your job that add up to major time include printed circuit boards, backplanes, card cage design and fabrication, and systems packaging. Our help here can put you ahead of the crowd and let you spend that time on your real business.



Characterized models speed simulation; 2D and 3D CAD files simplify board layout and package design. (Models available on CD-ROM.)


Packaging is a crucial element of systems performance, and our modelling, simulation, power and thermal analyses, and noise budget balancing are crucial for packaging success. When you're ready, our manufacturing experience—from clean room to global shop floor—can pay off at any level you choose: board, sub-system, or finished product.

Look to AMP for the help you need to meet tight production schedules and narrow market windows, and the resources you want to move into next-genera-



Complete system architectures include packaging systems, power supplies, and I/O. Developed and manufactured to your requirements.

tion technologies. Our knowledge base in fast boards, high-speed cable assemblies, surface mount technology, and related subjects can simplify the transition to faster logic and smaller form factors.




Custom high-performance cable assemblies represent value-added capabilities in advanced design technologies.

THE BIG PICTURE: FROM CONCEPT TO PRODUCTION.

Your program, and your problems, are unique. Your program is yours alone, but your problems don't have to be. We can help in any of the areas we've listed here.

Or let us put it all together for you, and move you from start to finish. Our specialists can bring our services to you anywhere in the world, and coordinate the AMP resources you need to bring your idea into production. Cost-effectively. And on time.



Printed circuit boards using advanced materials/technologies are the result of AMP alliances and advanced manufacturing facilities.

For more information, call our Product Information Center at 1-800-522-6752 (fax 717-986-7575). AMP Incorporated, Harrisburg, PA 17105-3608. In Canada call 416-475-6222.

THIS IS AMP TODAY.

AMP

continued from page 92

programs, iBrainmaker and DynaMind. The iBrainmaker program, developed by California Scientific Software (Grass Valley, CA), lets you simulate the network learning process through back-propagation techniques. In back propagation, you present the network model with a data set representing the application problem. Through simulation, iBrainmaker then trains the network to produce a desired response to specific inputs by assigning weights to each of the chip's analog storage elements. Once the network has been trained to solve the application problem, the weights are downloaded or programmed into the ETANN device.

The DynaMind simulation software, developed by NeuroDynamX (South Pasadena, CA), lets you simulate back-propagation learning, but also performs chip-in-loop learning. This technique optimizes the performance of the network by replacing the software simulation of the chip's performance characteristics and specifications with an actual device. The neural network can then "learn around" any minor processing variations occurring in individual ETANN chips.

Both iBrainmaker and DynaMind can also be used independently of the Intel development system to create neural network applications.

Fastest-learning chip

A more recent neural network IC is the RN-200, a 256-synapse (16 synapses \times 16 neurons) device that Ricoh (Tokyo, Japan) claims is the world's fastest-learning chip of its kind. The device boasts a front-end process of three billion connections per second and a learning speed of 1.5 billion connection updates per second (CUPS) when running at 12 MHz. In Tokyo, Ricoh demonstrated a desktop neural computer system that requires no software. Based on the first-generation RN-100 chip (a one-neuron device with eight synapses), the neurocomputer has a processing speed as high as 128 million neuron connections per second—a figure that Ricoh says will increase when the system incorporates the new RN-200.

One of the first neural network ASICs comes from Neural Semiconductor (Carlsbad, CA). Its CNU3232 has 32 inputs, 1,024 synaptic


Guns and Granny Smiths

Original and Mean Filtered Images

Original Image




Mean Filtered Image




Apple Inspection Example


Step 1
Collect Apple
Images



Step 2
Local Defect
Analysis



Defect Color
Codes



Step 3
Final Defect
Identification



In these examples of pattern recognition from HNC, the neural network's ability to learn by association shows the technology to be promising for everything from fruit grading to target recognition for weapons systems.

The apple grading system captures images and feeds them to a neural network, which eventually learns how to determine which characteristics affect an apple's quality. Once the network is trained, the system can classify an apple by comparing its traits to the network's learned base.

"The tank recognition system," says Ted Crooks, director of customer services at HNC, "was actually an experiment to prove that neural networks could be used in defense systems. Critics of neural networks cited an application where a trained neural system picked out tanks flawlessly until it was discovered that all pictures of tanks were taken in bright light and those without tanks were taken in shade or at night. As soon as those clues were taken away, the network failed. We proved that a neural network could indeed be trained to differentiate a tank's shape from other objects, and that experiment led to a real application."

In this picture, a filtered image of two tanks is outlined so that the network can learn to differentiate them from other objects.

For this Special Report on *Future Computing*, Computer Design interviewed Carver Mead, an expert on the subject of neural computing. Professor Mead is the Gordon and Betty Moore Professor of Computer Science at the California Institute of Technology, where he has taught for 20 years. He's also a co-founder and chairman of Synaptics, a company that develops neural network technology. Mead has pioneered in many areas of electronics, from the invention of the MESFET to silicon compilers and, recently, VLSI analog neural systems.

Computer Design: After many years of use in academic circles, neural computing now appears poised to move into practical commercial embedded applications. What's taken so long for this to happen?

Mead: Your question reminds me of how people used to talk about parallel architectures in computers years ago. People said: "We do things in a top-down way." And I'd ask: "Do you really know what the 'top' is?" I'd get these blank looks from people. But today, 15 years later, people are still sorting out what the top is in parallel computing. "Top-down" assumes you know everything in the beginning. That's not a very realistic view of the world. You don't know what's what in the beginning, so you have to evolve your understanding along with the application.

In the neural network business that top-down approach has translated into some rather abortive attempts to make general-purpose neural network chips. Those chips haven't worked well because no one knows what architecture is right for any real application. There are a lot of interesting simulations done in research circles, but none of those are applications that have to work in real time, under real-world conditions.

So then, rather than trying to generalize from basically no information—

which is what we're faced with today—it makes more sense to do specific applications. Until you've done that, there's really no royal road to the top to see what is the general case. General cases don't come that way. Those come hard-fought after years of working with specific cases. That's why I think it's important for designers to solve real problems all the way out. Then we'll begin to accumulate data that we can generalize from.

CD: How will future advances in VLSI process technology influence the capabilities of neural network chips? Do you see any potential roadblocks?

Mead: Silicon process technology is very relevant to neural computing. A lot of people have tried to invent brand-new technologies to do neural networks. But it's important to remember that we are riding on the coattails of a silicon technology that's highly evolved. Hundreds of billions of dollars have gone into this most advanced technology that civilization has ever known. And to think that you're going to start from scratch with some other technology and do as well is pretty silly. At Synaptics

we've been developing an adaptive analog technology that takes advantage of all the capabilities of process technology. As process technology evolves, it's immediately applicable to this adaptive analog approach. That's not true of other approaches to neural networks.

■ Transistors as analog devices

CD: As I understand it, your neural net chip uses the transistors in digital semiconductors in an analog way.

Mead: Yes. Transistors are analog devices. Let them be what they are. Digital IC designers have had to work so hard to turn them into 1s and 0s that they lose all that beautiful analog capability. But if you let them exhibit that capability, then you get to save a factor of 10,000 in power consumption, as we did in our chip. It's really remarkable.

The inputs and the intermediate sig-

nals are inherently analog signals—they're typically faked out by binary numbers right now in a computer simulation, but that's not the effective way to use them. The effective way is for them to evolve in real time as analog signals. Digital computers not only turn signals into digital values, but they also use discrete time. Those discrete time stamps actually destroy information by aliasing. And because a neural network is nonlinear, there's no theory that tells you how much information you've lost. Using transistors in the continuous (or analog) domain gets rid of those problems. We get back so much for using transistors in an analog way.


Fortunately, the transistors in semiconductors never did know that they were supposed to be digital. They're inherently analog by nature. That means you can use them that way, and the only thing you have to be careful of is that not all the transistors are exactly the same. That's why the adaptive part is so important. Because, not only do you let the neural networks learn from their environment, but they also adapt to changes in the environment. And one of the important changes in the environment is the difference among the transistors. If you do it right, the same paradigm that allows you to adapt to the outside world lets you adapt to differences among the transistors that are on the chip. That lets you use a commercial fab technology to build sophisticated neural net analog processors.

So far, we haven't seen any technological roadblocks in this path. In fact, we're seeing that as the technologies are evolving for digital use, they're actually evolving capabilities that we use in an analog way to make them even more effective.

■ The analog-digital continuum

CD: It seems clear that neural computing is not destined to replace traditional logical computing, by any means. It may even open up more opportunities for logical computing. How do you see the situation?

Mead: The real world tends to present you with continuous values: the intensity of a pixel or the value of a waveform, to take two examples. And the digital world deals with discrete symbols: the letters of the alphabet, for example. At some point, the continuous



stuff gets translated into discrete symbols. In the most general terms, this translation process is called classification.

When I started in electronics, almost the whole system was analog. And the classifier was a relay. Then you had a contact closure and that was your digital output. It turned on a light or a heater element in a furnace or something. Now we've gone hard over the other way to where we now have, in today's world, a little bit of antialias filtering and a multiplexer on the analog side. Then there's an analog-to-digital converter, which you could think of as the most brain-dead classifier you can imagine. Because all it does is take analog values and convert them into binary numbers which represent voltage changes. The computer is expected to manipulate these values as if they were numbers and eventually simulate a classification scheme that outputs discrete symbols.

The trend now is toward a much more balanced view of that picture. The analog preprocessing gets the data into a form where it's readily classifiable into appropriate higher-level symbols. If you were doing speech, for example, you'd want to classify the data into phonemes. Our I-1000 chip classifies images directly into character codes. A character code is a lot more meaningful to the computer than the analog values of the pixels. Computers are the way to handle discrete symbols, but they ought to be appropriate discrete symbols.

So we're seeing a return to letting the different technologies do what they're good at. You use an adaptive analog system to do all the preprocessing. You use an analog classifier to decide what the best classification is. That has a digital output which goes into a digital system.

Many of the neural network chips available today are aimed at that classifier job. But they're leaving out all the analog preprocessing that makes the data fit to be classified.

So, I believe that what we're headed for as a field is getting the picture to be more balanced, so the stuff at the analog end of the spectrum is done by adaptive analog technology and the stuff on the digital end is done by digital computers (as it is today). And the classification in the middle, between the two, is done at a level which makes sense for the problem.

■ The classification problem

CD: I guess you can shoot yourself in the foot if you make the classification problem too complex?


Mead: If you make that too big a problem, then it simply fails a lot. So you want to do those things at an appropriate level. And that's just plain good engineering. So we're not talking about replacing digital computers. It's just a matter of recognizing which technology is most natural for each situation.

CD: Today a lot of neural networks are simulated in software on large, powerful computers.

Mead: Right. It takes too much digital computing for the size of the problem. That limits the applications to the very few, where time isn't critical and you've got the processing power of a Cray computer around. It's only by getting this balance in the technologies that we're going to see very widespread use of the neural paradigm for real systems.

CD: Where do you see neural computing five or 10 years from now?

Mead: If you look at the continuum between analog and digital, we've gone all the way from 90 percent of the system being on the analog side to 90 percent on the digital side. We're headed back toward a balance of about half and half. Over the next few years, you'll see the analog side growing for applications which have to interface to the real world. There's no doubt in my mind that's where neural computing and computing in general are headed.



activation functions. The one-byte digital inputs and outputs and the weight-storage SRAM are all accessed through an 8-bit I/O bus. Unlike the Intel chip, which uses analog elements, the CNU3232 is a purely digital device that's targeted at embedded system applications.

"We refer to ourselves as a neural ASIC company," says Robert Bagby, president of Neural Semiconductor, "because we expect our customers to build neural network ICs of various sizes, topologies, precisions, and activation functions using our basic architecture. Neural networks are really multiple layers of nonlinear matrix multipliers. We build discrete circuitry for each and every neural multiplier, and place SRAM adjacent to that multiplier to store neural weight values. We also have a neural summation function built into the chip, so you have fully parallel neurons and fully parallel synapses or weights. Because our architecture is purely digital, designs based on it can be manufactured with standard processes for low-cost, high-volume implementations."

The first neural network IC to find its way into a commercially available product comes from Synaptics. The chip, designed for optical character recognition (OCR), hosts an analog sensing array, two neural networks and a digital controller on a single device. The chip is at the heart of a check reader being sold by Verifone (Redwood City, CA). "Until now, if you wanted to perform high-speed OCR, you were limited by the bandwidth between the sensor and the rest of the circuitry," says Synaptics' Bisset. "For most applications using a TV camera, that rate was just 30 images per second. By putting the sensor on the same chip with the classification circuitry, we can do the same task thousands of times per second."

■ Other solutions

Not all hardware solutions for neural network applications are based on neural-specific silicon. There are systems that use standard digital components for neural computing, as well as for non-neural applications such as Fourier or Gabor transforms. The CNAPS-1064 from Adaptive Solutions (Beaverton, OR) is such a hybrid—an array of 64 processing nodes (PNs), each with its own 4 kbytes of on-chip SRAM. A PN

DSP and neural nets team up for medical research


Any biomedical technician can probably do a better job than a computer of visually identifying reactions in blood cells. But when researchers at the University of California (Davis, CA) needed to classify 72,000 blood-cell reactions each day, and do so accurately and consistently, they moved to an automated approach that combines digital signal processing (DSP) with a neural network simulated in software.

At first glance, the application seemed like a simple image classification problem. Blood-cell images are captured using a high-resolution CCD (charge-coupled device) camera attached to a microscope. Each image is about a millimeter and a half in diameter. That area contains about 100,000 red blood cells. To induce a reaction, antibodies are suspended with the blood cells. After a period of incubation, a reaction may or may not occur. The photos show four degrees of reaction into which the neural network classifies images.

There's added complexity, however, because no two reactions are exactly the same. The network also has to account for rotated views of the exact same image. For example, if a human were looking at the image and felt he or she could better identify the reaction from another angle, the image could simply be rotated. For the computer this would be very complicated. You'd have to show the network how the image looked rotated 5 degrees, and that means the network would have to analyze 72 different versions of the same image. In back-propagation neural nets, the difficulty of training increases proportionally to the number of training examples raised to the third power.

■ Preprocessing needed

"In neural networks, the more work you can do before you hand off the project to the neural net, the better off you are," explains Wasyl Malyj, associate development engineer at uc Davis. With this in mind a preprocessing step was included to extract from the blood-cell im-



The neural network classifies blood-cell reactions into one of four types. The image in upper lefthand corner shows "no reaction" or a class 0 reaction. The image to the upper right shows some traces of clumping indicating a "weak reaction," or class 1. The lower left image shows even more clumping, indicating a "definite reaction," class 2. And finally, to the lower right the image indicates a "complete reaction," class 3, has occurred. Because these images are so complex (over a quarter of a million pixels), it's important to extract only the most relevant pixel data using DSP techniques. Otherwise, the neural network would require an unrealistic amount of computing power.

age only its most critical data.

To accomplish this, a two-dimensional Fast Fourier Transform (FFT) is performed on the image's pixel data, converting it into the frequency domain. This produces a compact feature vector. Sampling algorithms are applied to these vectors to extract useful information. The goal of the preprocessing is to take a very complex image with upwards of 512 x 512 (or over 1/4-million) pixels and extract from that a few hundred bytes of data. "The preprocessing reduced the amount of stuff that the neural net didn't need to learn, simplifying its structure, its training, and making it possible to implement the neural net with today's technology," says Malyj.

After the image is compressed into a complex feature vector, the vector is fed into the neural net. Cycling this information through the net lets it adjust its connection strengths, and in this way it

"learns" to associate particular spectral patterns with particular reactions.

Because of the preprocessing, the input stage of the neural net is typically 128 neurons. To implement the network, the uc Davis researchers developed in software a custom-written back-propagation simulator capable of building nets with three or four layers. The code was written to run on a 486 working in conjunction with a Motorola 96002 floating-point digital signal processor (DSP).

The input layer typically has 128 neurons. The hidden, or middle, layers usually have between 12 and 64 neurons, while the output layer consists of 8 to 20 neurons. Four of the output neurons represent the discrete reaction classes, an additional output provides a confidence metric and the remaining outputs code for a variety of possible error conditions, such as cracks in the plastic tray, bubbles in the sample, shadows cast by bubbles in the mineral oil covering the specimen, lack of blood or reagent in the reaction well, and proper focus.

■ Training the net

According to Malyj, it took only about six hours to train the neural net. The DSP hardware helped

boost its speed. A training set of about 800 images was shown to the neural net, along with the classification under which each image belongs. Once the neural net was trained, technicians began to feed it images that it had never seen before for classification.

The network can be adjusted for various performance levels. "If we tell the neural net to classify absolutely everything, it's accurate to percentages from the high 80s to the low 90s," says Malyj. "That's not as good as a human. But if we tell the net to classify only those images about which it's 'confident,' and to flag the others for us to take a look at, then it classifies about 85 percent of the images at better than 99 percent accuracy."

Researchers can then take the remaining 15 percent of the images and, after they've been scored by a human, use them to retrain the network.

resembles a simple digital signal processor (DSP), and can be programmed for a variety of applications. At 25 MHz, the device can compute 3.2 billion multiply-adds per second. The chip falls somewhere between PC-based software solutions for neural network applications and silicon that's targeted at those applications.

"Given the performance of today's PCs and workstations, you can emulate a lot of neural network applications in software alone," says Dan Hammerstrom, founder and chief technical officer at Adaptive Solutions. "But from a silicon standpoint, there are two approaches you can take. One is a device like ours, which can be customized for a certain domain of applications but is flexible enough to be programmable. The other is a chip like Synaptics', which is designed for a specific task and which may have a price/performance advantage for that application. Both have their place. For complex tasks such as Kanji character recognition, it's simply not feasible to build a custom chip. That's where flexibility pays off."

■ Software-only solutions

The remaining category of neural network-specific products is software-only solutions, which emulate neural networks on workstations, PCs and mainframes. These programs use either standard microprocessors or DSP chips to perform neural networking tasks, and are available from dozens of companies. The applications are as diverse as the companies that offer them, but most products are being used for pattern recognition, financial analysis and defense-related projects.

In essence, most neural network tasks are based on some form of pattern recognition. Sometimes the sought-after pattern is a visual one—for example, a system that sorts good fruit from bad on a conveyor belt. In this application, the network must be trained to look for a vague trait such as "quality." "The customer who wanted to grade fruit


had to train the network to recognize color patterns that distinguished good apples from bad," says Ted Crooks, director of customer services at HNC (San Diego, CA), a neural network software vendor. "By repeatedly presenting data to the network, he trained it to discern

cell structure, some physicians are using neural networks to help them with diagnosis and prognosis. "One of our customers is a neurosurgeon who's using neural networks to predict potential IQ loss after brain surgery," says Jim Blodgett, director of marketing for California Scientific Software. "Traditional statistical methods have been used in the past, but they lose precision in the middle of the bell curve. At each end, either with slight damage or severe damage, the statistics are fine. But in the middle of the curve there can be large variations. Neural networks can use factors such as the severity of an injury or a patient's medical history to make predictions of an operation's outcome."

■ Financial uses

While neural network research in medicine makes for dramatic reading, there are other applications, particularly in the realm of finance, that might have even greater ramifications. Banks are relying on neural networks to do everything from predicting loan eligibility to spotting credit-card fraud. In the case of loan eligibility, the network is taught to examine the factors that would make a good loan applicant, based on profiles of good and bad loan recipients. The usual criteria are included in the data, such as income, time on the job, credit history, and so on, but neural networks look for unusual patterns or relationships which might escape a human, particularly a human who looks at dozens of applications a day. In the case of credit-card fraud, the network might look for spending patterns which are unusual, such as someone who purchases a wide-screen TV, a trip to Europe and a gourmet meal in a 24-hour period.

The common denominator in these applications is training the network to look for subtle shifts in patterns which are crucial to making a final decision. In the case of stock market predictions, one analyst said that a day on Wall Street is akin to being hit with a dozen data



"Not all neural networks are created equal," says Steve Ward, president of Ward Systems Group. "A traditional feed-forward network produces one and only one set of outputs from a given set of inputs after it's trained. A recurrent network, on the other hand, may produce many different sets of outputs from one given input set, depending on when in time the input set is presented to the network. Recurrent nets not only learn individual input patterns, but learn sequences in those input patterns as well."

the important relationships that define good quality—for example, 'this is premium,' 'this is grade six,' 'this is rotten,' and so on. He didn't define what made up these classifications; he just gave examples of what characteristics are needed to place a piece of fruit there."

Visual pattern recognition is also being applied to medical research. Some hospitals are using a neural network to sift through hundreds of slides to detect anomalies in blood cells or tissue samples. Early results from these applications show that neural networks achieve a surprising level of accuracy when they're compared to a human performing the same task. As with most applications where a computer equals or bests a human being at a task, fatigue is the deciding factor. Although human expertise may hold the upper hand in picking a cancerous cell out of thousands of normal cells, human fatigue can cancel out some of that expertise, and the capabilities of human versus neural network begin to equalize.

In addition to recognizing flaws in

firehoses at once. The trick is to ascertain which droplets of data in that enormous stream are the ones that will affect the market.

"Naturally, the key here is giving the system data," says Steve Ward, president and technical director of Ward Systems (Frederick, MD). "One of our customers, an investment advisory firm, is using our NeuroShell software to create a Standard and Poors stock market prediction system. The client uses technical indicators to predict short-term changes, and a combination of technical and fundamental indicators to predict long-term changes. The database includes 12,000 companies

and looks for fundamental data patterns in areas such as growth and cash flow yield, as well as cross-sectional analysis or valuation change to compute expected return. Early experiments using a model stock portfolio for the years 1987 through 1990 have yielded impressive results in predicting which stocks produce returns better or worse than expected."

■ The lure of networks

Obviously, developing neural networks that can accurately predict what was once thought to be unpredictable is a tantalizing prospect. At present, people are using them to

guess the outcome of everything from the effects of a war on oil prices to the outcome of a horse race. It's true that a lot of these stories sound like hype, especially when they're playing to an audience of computer professionals who've seen their share of flash-in-the-pan technologies over the years. Still, there is an element of mystery to many neural network applications.

"We think there's a lot more going on, particularly in financial circles, than people are letting on," says Adaptive Solutions' Hammerstrom. "After all, if you had the inside story on the stock market,

Neural nets give tin ears a good name

Many applications that are suited for neural computing are tasks at which humans are better. But computers have an advantage over humans. They don't get tired or bored—even after several hours of repetitious work. With this in mind, engineers at CTS (Matamoros, Mexico) made use of a neural network in their loudspeaker manufacturing process.

At its plant, CTS manufactures several million loudspeakers per year. To ensure the quality of the units, a final inspection was performed by a trained operator, skilled at identifying audio defects.

This method had some disadvantages. An operator was required to listen to 2,000 or 3,000 speakers in a single day. Since the final test is very subjective, over the course of a day the operator's fatigue level, emotional state and stress level affected the evaluation. And, because the test is subjective, pass/fail criteria varied from person to person. As a result, many perfectly good speakers were rejected, reducing the factory's yields.

To solve this problem, CTS wanted to remove the subjectivity from the testing process—or at least to make testing consistent from one production run to the next. At first the company used PC-based audio test equipment that was sensitive enough to measure distortions caused by speaker defects.

But this alone wasn't enough. Testing also had to

classify the units as good or bad. At first, statistical pass/fail limits were enforced. These turned out to be very unforgiving and depended too much on consistent equipment and fixture setup. In addition, the equipment didn't distinguish among the various types of defects.

It was at that point that Rick Bono, a design engineer at CTS, decided to try using a neural network to classify the results of production testing. A back-propagation neural net was established, consisting of 10 input nodes, 18 hidden-layer nodes and 4 output nodes.


To begin with, a piece of test equipment measures the loudspeakers at different frequencies. The distortion values at these frequencies are the inputs to the neural net. The four possible outputs identify good speakers and three classes of defective units.

To train the neural net, CTS ran over 200 units, both good and bad, through the process, while manually classifying them into one of the four output classes. This required about 20 minutes of computer time on a 25-MHz 486. Once the network was trained, Bono integrated it into the test software using the NeuroShell run-time code generator from Ward Systems Group (Frederick, MD). Network values were developed for each speaker model.

CTS has been using the neural-based system for a year and a half and now runs all its products through this testing process. Over this period of time, the company has seen significant improvements in production yields, thanks to the neural net's ability to classify consistently.

Bono is now considering developing a second-generation system which would incorporate new cases as it works, train invisibly and then implement the new network. This would require a hardware-based implementation of back-propagation, says Bono.

Other upgrades might include adding more output nodes to the neural net. "We've seen different cases pop up—different defect cases that weren't included in the original training," says Bono. "They have distinct patterns, and the net can be trained to recognize them."



DEVELOPMENT TOOL VENDORS

List courtesy of Martin Middlewood and Tom Schwartz.

★ Adaptive Solutions

CNAPS: Development environment includes CNAPS assembler and library of neural network algorithms.

(503) 690-1236 / Circle 366

AI Ware

N-NET EX: User and program interfaces, functional link net architecture, associative recall. Supervised and unsupervised learning.

(216) 421-2380 / Circle 367

AND America

HNNet: Neural-based development system using digital holography principles. Transputer- and PC/Windows-based versions.

(416) 569-0897 (Canada)
Circle 368

Applied Cognetics

WinBrain: Develops back-propagation networks. Incorporates multiple transformational models.

(212) 969-8769 / Circle 369

★ California Scientific Software

Brainmaker: Basic version supports up to 512 input neurons, up to six hidden layers. Tutorial and eight sample networks. Imports Excel, Lotus, dBase, binary, and ASCII files. Print/edit neuron matrices.

(800) 284-8112 / Circle 370

EPIC Systems Group

Neuralyst: Integrates neural networks with Excel spreadsheets. Includes macro library for investment analysis.

(818) 564-0383 / Circle 371

★ HNC

ExploreNet 3000: Windows-based application software program for developing and implementing neural network solutions without programming. Database mining program available also.

(619) 546-8877 / Circle 372

Hyperlogic

Owl Neural Network Library: Twenty-four functions for accessing networks supplied as C library. Twenty types of neural networks.

(619) 746-2765 / Circle 373

ImageSoft

ExperNet: Object-oriented tool for creating Windows-based neural networks and knowledge applications.

(800) 245-8840 / Circle 374

Inductive Solutions

NNetSheet: Supports nine algorithms for supervised and unsupervised training. Train network can be ported to a spreadsheet.

(212) 945-0630 / Circle 375

Mathworks

Neural Network Toolbox: Includes learning rules, transfer functions and training and design procedures for implementing neural networks.

(508) 653-1415 / Circle 376

Martingale Research

SYSPRO: FORTRAN-based neural network simulation and prototyping tool.

(214) 422-4570 / Circle 377

Neural Computer Sciences

NeuralDesk: Supports many algorithms. Manual and automatic training of neural networks.

44-703-667775 (UK)
Circle 378

Neural Systems

Genesis: Development environment for interfacing neural networks to application software.

(604) 263-3667 (Canada)
Circle 379

★ NeuralWare

NeuralWorks: Neural network chip development, open architecture, 8-k back-propagation, makes network types from libraries and creates diagnostic tools.

(412) 787-8222 / Circle 380

Neurix

MacBrain: Flexible neural connections, activation rules, 3-D graphs, interactive

modeling, visual macro language.

(617) 426-5096 / Circle 381

★ NeuroDynamX

DynaMind: Train networks on Intel's 80170NX ETANN and Intel multichip board. Can read and store network trained in emulation mode and download weights to chip.

(800) 747-3531 / Circle 382

NeuroSym

Neural CASE: Supports four network paradigms: BPN, CPN, RN, and SOM.

(713) 523-5777 / Circle 383

Peak Software

Autonet: Constructs networks from training data sets consisting of input variables and expected results. Networks may also be created from command line.

(612) 854-0228 / Circle 384

SAIC Artificial Neural Systems

Delta ANSpec: Language for defining and implementing parallel distributed processing systems.

(619) 546-6005 / Circle 385

Software Bytes

ET 2.0: Simulates text, graphics and Windows. Back-error propagation neural networks with Borland C/C++ source code, ET Graphics and Windows slide neural networks on equivalent VGA screens.

(800) 521-4119 / Circle 386

Software Frontiers

Neural Network Toolkit: Development software for neural network applications. C source code included.

(800) 475-9082 / Circle 387

Talon Development

Brain: Lotus 1-2-3 add-in for neural net development.

(414) 962-7246 / Circle 388

★ Ward Systems Group

NeuroShell: Shell program imports ASCII spreadsheet problem files. Example included. Windows version can build up to 128 interacting networks. Supports dBase. Includes run-time option.

(301) 662-7950
Circle 389

CHIP AND HARDWARE VENDORS

★ Adaptive Solutions

CNAPS System: A 5-billion-connection-per-second neurocomputer. The system has a back-propagation learning rate of 1 billion connection updates per second.

(503) 690-1236 / Circle 390

American NeuralLogix

NLX420: Neural processor slice. A digital chip designed for real-time neural network systems. This 20-MHz device contains 16 processing elements, and can have up to 64,000 16-bit synaptic inputs.

(407) 322-5608 / Circle 391

★ Intel

80170NX ETANN: An electrically trainable analog neural network chip. One chip can perform over 2 billion multiply-accumulate operations per second.

(408) 765-9235 / Circle 392

★ Neural Semiconductor

CNU3232: Digital neural net chip implements a single-layer network of 32 inputs and 32 nodes, capable of processing 100,000 patterns per second.

(619) 931-7600 / Circle 393

★ Ricoh

RN-200: A neural net chip that implements a 256-synapse neural net composed of 16 synapses by 16 neurons. The chip is a 200,000-gate array built on 0.8-μm CMOS. It has a forward process of 3 billion connections per second and a learning speed of 1.5 billion connections per second.

(408) 432-8800 / Circle 394

★ Synaptics

I-1000: Analog neural network chip designed specifically for reading checks. A neural network-based image sensor reads the image and a neural network trained to recognize the characters on a check interprets them.

(408) 434-0110 / Circle 395

would you boast about it, or quietly use it to get rich?"

Perhaps the only area where more mystery prevails than on Wall Street is in defense-related neural network applications. Most of the people involved in such projects are understandably reluctant to discuss the details of their activities, but it's clear that neural networks are being used for such things as missile guidance systems.

"We became interested in neural networks about six or seven years ago, when other artificial intelligence solutions proved to be too slow for our applications," says Dr. David Andes, research fellow at the Naval Air Warfare Center (China Lake, CA). "As you can imagine, the life of a missile is very short—often only a few seconds. The amount of room in which you have to store a guidance system and its power source is also pretty limited. After all, the purpose of the device is to deliver explosives, not electronics. We got into neural networks because biological brains do the type of computing that we need, and they do it very fast and in a very small area. We're trying to build a guidance system that can hit a target without needing human intervention."

Naturally, for these applications a neural network must be trained to recognize a target in the confusion of battle, something that heat-seeking devices find problematic. But trying to give a neural network enough data to find targets in rapidly changing battle situations is a daunting task.

"Neural networks are notorious for picking up on things that you don't want them to," Andes cautions. "We heard about one application where a neural network was picking out the enemy from a visual field with perfect accuracy. Naturally, everyone got suspicious and investigated more closely. It turned out that the system produced a 60-cycle hum whenever the picture of the bad guy was shown, so the network just incorporated that into the equa-

tion—if hum, bad guy; if no hum, good guy. They took away the hum and the network was lost."

Neural network's impact

When the breadth of neural networking applications is examined, it's clear that if they're refined, they will affect our lives as much as any technological breakthrough of the twentieth century. But trying to find out where to separate fact from fiction is difficult, particularly when those who are really successful with neural networks are reluctant to di-

ing this technology to analyze failures at its fabrication lines. Location information is reported in aeronautics-style polar coordinates. For example, a defective side of a 6-in. wafer may be at 6 o'clock—like a pilot reporting an enemy's position in the sky.

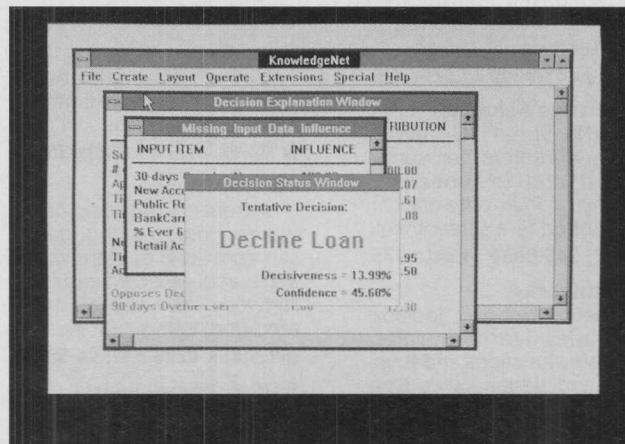
Neural networks could also conceivably guide placement and routing algorithms for chip and printed circuit board design, in essence making decisions intuitively, like experienced engineers, rather than through slavish adherence to algorithms. The potential is only limited by your imagination and by the amount of data that you have.

"When we qualify potential customers' applications, we go by how much organized data we have on them, not on the size of their checkbooks," says HNC's Ted Crooks. "There are plenty of companies out there who have all this data lying around. With a little ingenuity we can turn a large database that used to be a nuisance into an asset. But without that data we don't even bother. All we'd get would be an aggravated customer."

It's true that we are a society awash in data. The government and the banking industry alone probably have enough

statistics about most of us to predict our voting habits, buying patterns and loan eligibility.

If all of this sends an Orwellian chill up your spine, you're not alone. Although it's far too soon to predict whether neural networks are another headline-grabbing AI story or the beginning of a computer revolution, one thing is clear. The photographs that accompany this article are taken from real applications. Somewhere a computer might be grading the apples that you eat. It may be deciding whether to give you a loan or recommend you for a job interview. And because most of you reading this article are familiar with what computers can and can't do, you're either smiling right now—or feeling a little queasy. ■



Banks are using neural network software to verify a person's loan eligibility. In this application from HNC, variables such as an applicant's time at a fixed address, time at a particular job, previous loan delinquency, and savings account data are passed through the network for evaluation. The network builds a profile of the loan applicant and makes a recommendation to the financial institution.

vulge too many details. And so the questions remain—can you use a neural network in your job, and how will neural networks ultimately affect your life?

As far as jobs are concerned, neural networks are best suited to analytical tasks that prove too complex or tiring for humans to perform accurately. And certainly because neural networks are based on computer technology, they will affect the electronics industry if they're widely embraced.

Technological applications of neural networks are, in fact, starting to see the light of day. Last February, for example, Intel announced a breakthrough in neural networking, the capability not only to identify patterns but also to read out their locations. The company is us-

A PennWell Publication

| ISSUE | SPECIAL REPORT | TECHNOLOGY FOCUS | PRODUCT FOCUS | DESIGN STRATEGIES | OEM INTEGRATION |
|---|--|---|---|--------------------|--|
| JANUARY | Verilog and VHDL | • DSP development tools • Disk drive controllers and associated ICs | C cross compilers | Networking | • Secondary & mezzanine buses • Backplanes & enclosures |
| FEBRUARY <i>Buscon West, EDAC/Euro ASIC</i> | Bridging the buses | • Software tool integration • ICs for desktop video | Flash EEPROMs | Process control | |
| MARCH <i>RISC '93, PCB Design</i> | Memory architectures | • Multithreaded operating systems • PC-Based CAE/CAD Tools | SCSI host adapter boards | Data acquisition | • Input devices • Printers & output devices |
| APRIL* | EMBEDDED COMPUTER CONFERENCE (ECC)** SHOWGUIDE | | | | |
| ECC | PC/AT architectures in embedded applications | • Developments in 3-V ICs • Evaluating simulation strategies | Device programmers | Peripherals | |
| MAY | SPECIAL REPORT ON FUTURE COMPUTING: Virtual Reality | | | | |
| <i>Comdex Spring, CICC</i> | New applications for DSP | • Benchmarking programmable devices • Bus standards | Video D-A converters | Portable computers | • Networking interfaces, standards & components |
| JUNE* <i>DAC</i> | High-level synthesis and architectural design | • Small form factor VME • Data compression standards and ICs | Ultra-fast SRAMs | Graphics | |
| JULY <i>Fuzzy Logic '93</i> | Advances in IC packaging | • Interfaces for DSP • Fuzzy/neural update | Real-time kernels and operating systems | Imaging | • Display devices & monitors |
| AUGUST* | Trade-offs in programmable devices architectures | • RISC in real time • Integrating CAE and CAD databases | Emulators | Robotics | |
| SEPTEMBER <i>EuroDAC, Buscon East Embedded Systems</i> | Software testing and quality | • Futurebus+ • Integrating testability into the ASIC design process | Low-power DRAMS | Instrumentation | • Power sources • Interconnects |
| OCTOBER* <i>Analog & Mixed-Signal</i> | ANALOG & MIXED-SIGNAL CONFERENCE SHOWGUIDE** | | | | |
| | Designing mixed digital/RF systems | • Hardware/software trade-offs in multiprocessing • Network interfaces and interface ICs | Logic analyzers | Simulation | |
| NOVEMBER | SPECIAL REPORT ON FUTURE COMPUTING: The merging of computers and communications | | | | |
| <i>Wescon, Comdex</i> | Designing for testability and manufacturability | • PLD design tools • Mezzanine buses | High-resolution A-D converters | Communications | • Mass storage (disk, tape, CD-ROM, memory cards) |
| DECEMBER* | Designing the next generation of portable computers | • Full-system simulation • FAX/modem ICs | VME CPU Boards | Supercomputers | |
| Starch Readership Readership Issue | | **Contact Computer Design for circulation, rates and closing dates | | | |

PENNWELL PUBLISHING COMPANY • ONE TECHNOLOGY PARK DRIVE • WESTFORD, MA 01886 • 508-692-0700

For more information, contact any of us: PUBLISHER: David L. Allen, 508-392-2111 • ASSOCIATE PUBLISHER: Tim L. Tobeck, 508-392-2116 • NATIONAL ACCOUNTS MANAGERS: • WESTFORD, MA, Tim Pritchard, 508-392-2217 • SPRING, TX, Eric Jeter, 713-353-0309 Fax 713-288-8350 • SAN JOSE, CA, Diane Palermo, 408-371-7551 Fax 408-371-6239 • STRATEGIC ACCOUNTS REPRESENTATIVES: • WESTFORD, MA, Sue Nawoichik, 508-392-2118 800-223-4259 • David Singer, 508-392-2109 800-223-4259 • INTERNATIONAL SALES OFFICES: • SURREY, ENGLAND, David Round, 81-686-7655 Fax 81-688-2134 • PARIS, FRANCE, Daniel R. Bernard, 1-39-14-6780 Fax 01-39-14-7014 • MUNICH, GERMANY, Johann Bylek, 089-903-8806 Fax 89-904-3526 • MILAN, ITALY, Luigi Rancati, 2-70-30-0088 Fax 2-70-30-0074 • TOKYO, JAPAN, Toshio Equis, 81-33-536-5404 Fax 03-536-5490 • ASIA, AUSTRALIA, NEW ZEALAND, Tom Gorman, 852-833-2181 Fax 834-5620



The path to top-down design can be anything but straight and narrow. And once you get off track, you can get into all sorts of trouble. The ASIC wilderness is littered with good point tools that somehow fail to provide an integrated solution.

COMPASS® Design Automation knows all about the rules of the road. That's because we've been down it more than any other EDA company. We've developed the

most integrated top-down ASIC solution in the industry, from first concept to first silicon. Supporting even the largest and most complex designs that require a million gates and more.

Simply put, we save you time. Starting with an easy-to-use integrated graphical HDL and synthesis environment through comprehensive cell-based and gate array libraries. And all the automated tools you need to move on down the design road.

So chart your course today. For more information, please circle the reader number, call **(800) 433-4880, ext. 7097** or fax your request to **(408) 434-7977**.

COMPASS 
Design Automation

1865 Lundy Avenue
San Jose, CA 95131
Tel: 800-433-4880
Fax: 408-434-7977

Caution: ASIC design may be hazardous to your health. Design delays due to poorly integrated tools may have a detrimental effect on your career. COMPASS assumes no responsibility for those who fail to follow our path to Silicon Success.


Analog designers still trail behind their digital counterparts

Still dealing with custom more than semicustom, analog designers plod along the trail without the help of synthesis, a standard analog HDL or test methodology. But the determined can still hit pay dirt with the tools at hand.

It may be small steps that users and vendors are taking to span the chasm dividing analog and digital design methodologies, but those willing to brave the pitfalls are excited with the results they're getting. At Siemens Gammasonics (Hoffman Estates, IL), where automated tools have finally replaced breadboards and scopes, for example, designers have discovered it takes just two or three hours to do with a simulator what had previously taken three or four days. Group leader Roger Arseneau reports first-pass success on a 10-bit D-A converter plus preamplifier that's going into a nuclear medicine camera.

Arseneau uses the Fastrack analog ASIC design system from Harris Semiconductor (Melbourne, FL). Fastrack tools form

Barbara Tuck, Senior Editor



SPECIAL REPORT: ANALOG DESIGN

the basis for the Analog Artist toolset from Cadence Design Systems (San Jose, CA), and fit into the Cadence framework. "They're complicated tools," says Arseneau, "but Harris gave us lots of support. There were usually two people manning the phone, so I had answers in three or four minutes."

Since he was using a new current-mode logic family developed for this particular ASIC project, Arseneau didn't have precharacterized cells to work with. "It was all custom de-

ably to a discrete version's drift of 200 to 400 μ V. The mixed-signal chip is on its way into production now.

At Boeing Semiconductor Design Engineering (Seattle, WA), advanced technology group leader Bryan Buchanan also relies on Harris' Fastrack. Buchanan uses the Cadence-based Harris toolset for the analog portion of the mixed-signal ASICs he's designing for the cabin management system of a Boeing 777. For mixed-mode simulation, however, Buchanan turns to LSim from

(AHDL)? "At Harris," he says, "designers create behavioral models in FORTRAN or C. It's not as elegant as an AHDL. We'd like to have a standard."

Interface models a problem

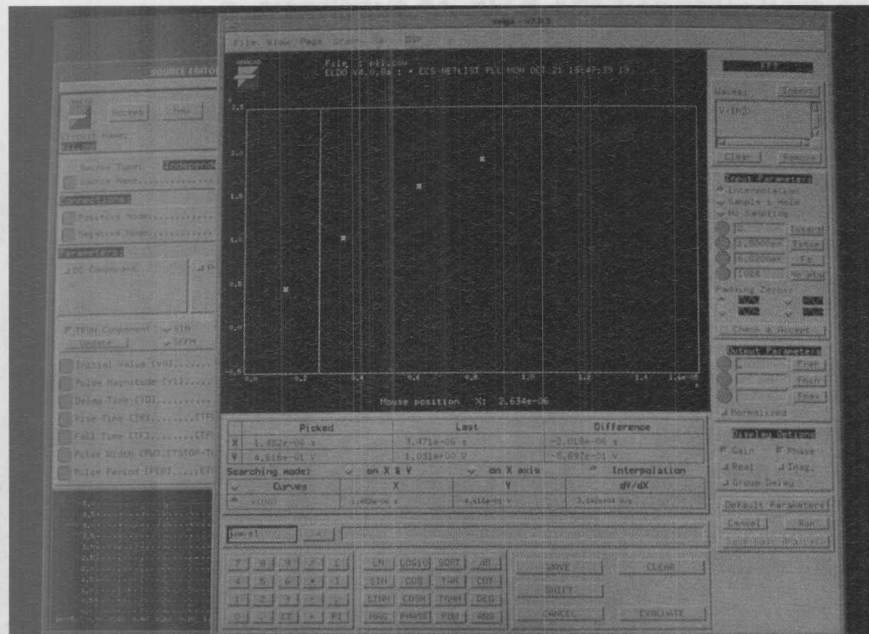
Coston adds that dissatisfaction with the interface models between Cadence's Spectre Advanced Circuit Simulator and Verilog is keeping Harris from using Spectre for mixed-mode, even though Cadence has been shipping an integrated Verilog/Spectre for several months now. Harris is still using Cadence Spice with Verilog for mixed-mode simulation, both in-house and in toolsets being shipped to users.

Among the issues that need to be addressed when building interface models, according to Coston, are loading, rise and fall times and forward/backward time-step control. Cadence is working on the interface models, and Harris intends to integrate Spectre into its design systems within a year. Nader Fathi, director of analog applications engineering at Cadence, reports that the company is also looking at a simulation backplane that would let more than two simulation engines run simultaneously. Today, Cadence uses the Interprocess Communication Protocol to let two simulators communicate with each other.

Although Harris doesn't yet use Spectre for mixed-mode simulation, designers at the company did use it in stand-alone mode to ensure the system-level functionality of its 10,500-transistor, 12-bit, BiCMOS sampling A-D converter, the HI-5800. While Cadence Spice worked well for smaller block designs in the HI-5800, it couldn't simulate the entire system or even get past the reading-the-circuit stage. Spectre read and simulated the HI-5800 cleanly in one sweep, however, avoiding many of the errors that occur when a circuit is partitioned into smaller components. If it weren't for Spectre, according to Harris, the production date of the HI-5800 might have been pushed back by several months.

From custom to semicustom

Designers at Monolith Technologies (Tucson, AZ), a joint venture of Tucson's Burr-Brown and Taiwan's Hualon Microelectronics, are counting on synthesis tools to speed them through a five-month design cycle. "We're counting heavily on all the tools we have to get our mixed-signal chip out in time," says Dr. John Man-



Having a user-friendly interface that lets you do all the common operations you're looking for when designing a circuit is one requirement for today's simulators. The Eldo simulator's XELGA waveform post-processor lets various mathematical and DSP functions be applied to signals saved during simulation. Here, Eldo plots the simulation result from a single signal where points are chosen on the resultant curve and on which mathematical computations are performed. This particular example is a fast Fourier transform.

signed. I had to simulate at the cell level," he says.

Arseneau simulated both analog and digital with Cadence Spice—until recently referred to as cdsSpice, an outgrowth of Harris' in-house Spice simulator. "Simulating the entire chip as analog did slow it down," he says.

The preamplifier was the most critical portion of the mixed-signal design, but Arseneau found he could do things to ensure temperature stability that had been impossible in the lab. "There was no way I could change the size of an emitter to see if I could get better temperature drift in the lab," he adds. A temperature drift of 30 μ V over a 40°C change did, in fact, compare favor-

Mentor Graphics (Wilsonville, OR) and the HSpice circuit simulator by Meta-Software (Campbell, CA).

Harris also offers an extension of its analog Fastrack toolset for mixed-signal design. Called Mixed-Signal Fastrack, it includes Cadence Spice with Verilog for mixed-mode simulation, as well as a charge-based, switched-capacitor simulator and a macromodel synthesis simulator. W. Terry Coston, director of design systems at Harris, says the company has been getting a lot of requests for Mixed-Signal Fastrack this past year, as more and more designers look for a mixed-signal solution. How does Harris get around the absence of a standard analog hardware description language

Put your system on silicon.

CB-C7 takes the sweat out of cell-based custom.

Our exciting new CB-C7 cell-based ASIC technology makes putting your system on silicon almost as easy as a day at the beach.

You get megafunctions, astonishing ECL-like speed and dramatically lower power consumption. And, migration from gate array designs is easy, because CB-C7 uses the same cell names and cell functions as our CMOS gate array families.

Use the right tool for the job.

NEC's OpenCAD® Design System integrates standard design tools with our proprietary tools, so you can get the job done at your worksite.

We support most leading vendors including Synopsys, Cadence, Vantage, Viewlogic and Mentor Graphics.

NEC reliability is your life preserver.

NEC gives you proven functions, proven testing and a reputation for quality and reliability known throughout the world. And now, our expanded Roseville, California facility brings high-volume submicron manufacturing even closer to you.

To learn more about CB-C7 cell-based custom design with embedded CPU cores, compilable RAM and ROM, megafunctions and more, call the toll-free number.

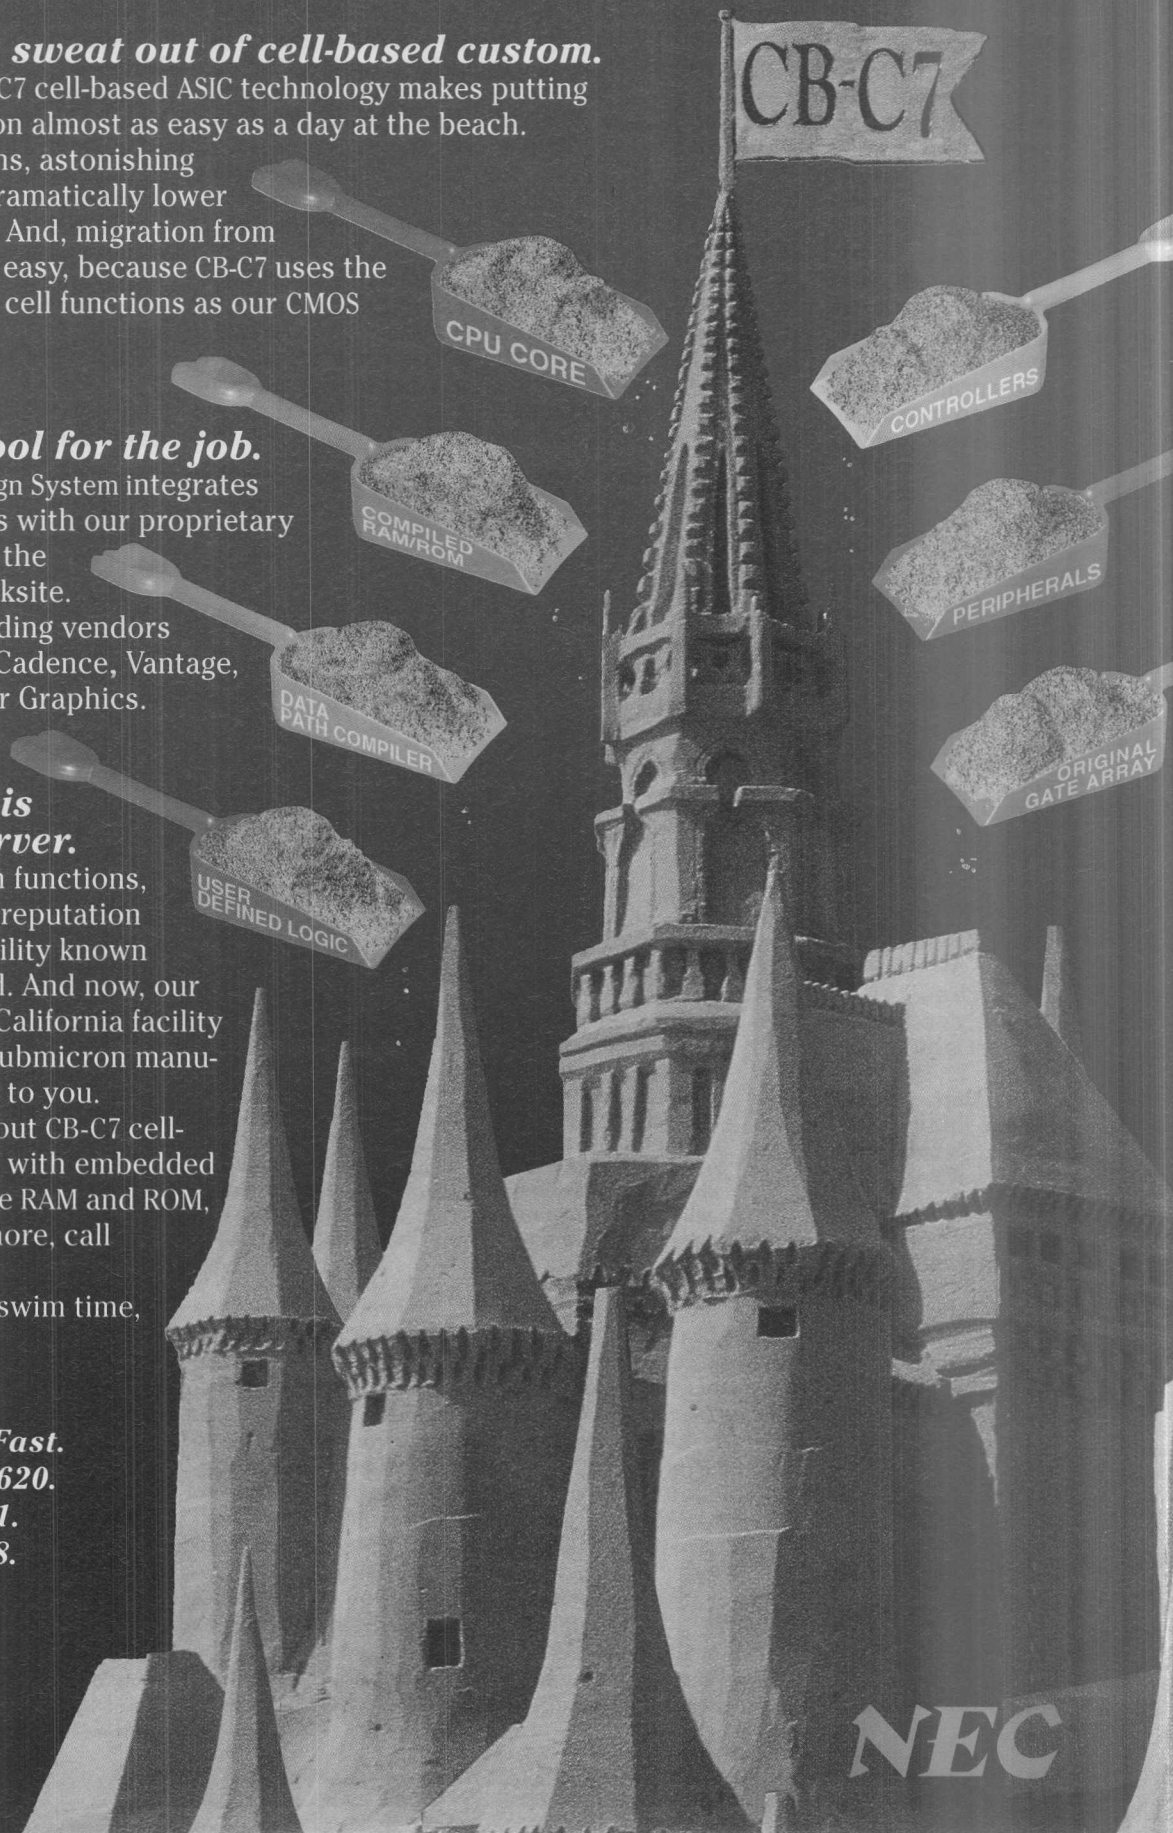
When it's sink or swim time, choose NEC.

***Free Information Fast.
Ask for Info Pack 620.
Call 1-800-632-3531.
FAX 1-800-729-9288.***

© 1992 NEC Electronics Inc.
401 Ellis St. P.O. Box 7241
Mountain View, CA 94039-7241
1-800-632-3531

OpenCAD is a registered trademark
of NEC Electronics Inc.

CIRCLE NO. 90



NEC

pany. Headed for the Japanese market, the multimedia chip combines 20-bit analog performance with high-speed digital signal processing. "The Japanese market works on a six-month cycle," adds Manning. "If we miss by a month, revenues for the product are cut in half."

Although Monolith designers built their own cell library for the multimedia product, making it a custom design, Manning says their approach was really akin to ASIC methodology, since they will use the same cell-based building blocks for future projects. AnaCAD tools, for analog and mixed-mode simulation, were complemented by tools from Comdisco Systems (Foster City, CA) for high-level DSP design and simulation, as well as logic simulation and layout tools from Compass Design Automation (San Jose, CA).

Monolith was the first U.S. company to purchase a license for the Eldo simulator from AnaCAD, now headquartered in Fremont, California, as well as in Ulm, Germany and Grenoble, France. "We have every tool that AnaCAD sells," says Manning. "They give us all the support we need—and very quickly. We give a call to the technology support person, and we get a response right away."

Eldo is the heart of a simulation system that includes analog behavioral (FAS), accelerated (XL) and mixed-signal (Fideldo) simulation—the latter combining Eldo and Fidel modules. The system also has modules for circuit optimization (OPT) and digital signal processing (DSP). Eldo, Eldo-XL and Eldo-FAS are time-step-driven simulators, whereas Fidel is event-driven. The synchronization of the time-step and event-driven simulators is handled by AnaCAD's Mixed-Signal Procedural Interface, a library of communication functions optimized for speed and accuracy, with both a voltage and a current interface.

■ Simulation speed improved

The recently released Version 4.1.0 of the Eldo analog and mixed-signal simulation system improves simula-

ing mixed analog and digital MOS devices, and it's also said to significantly upgrade Eldo's FAS analog behavioral modeling option. By improving the decoupling between the two solution algorithms Eldo uses in its divide-and-conquer approach to simulation, AnaCAD has given users the ability to separately control the time-step size and accuracy of each algorithm.

According to Peter Denyer, AnaCAD's vice-president for marketing, "The time-step size and accuracy needed for circuits operating in a digital mode is typically orders of magnitude greater than that required for circuits operating in an

gained by automatically translating the model into C source code.

Designed to be fast to simulate large analog/digital circuits, Eldo matches the speed of Spice for small circuits but is much faster for circuits with over 1,000 transistors. But it isn't just speed that sold Monolith's Manning on Eldo. "We don't care as much about speed as we do about accuracy," he says. "We can always plug in a faster CPU card or let a simulation run over the weekend."

■ Benchmark crucial


The pivotal test for Monolith engineers was a benchmark the company performed. Manning was intri-

guished that there was no clear winner among U.S. tool vendors participating in a benchmark study at the 1990 Bipolar Circuits and Technology meeting—the so-called BCTM benchmark. The benchmark involved simulating and predicting the conversion time for an all-bipolar 12-bit A-D converter. A prerelease version of Cadence's Spectre did predict the conversion time, but it had to use some strange techniques to do so, according to Manning.

The Monolith president went to the source of the circuit being simulated to do his own benchmark. "All the Spice simulators we used up until then were unsuccessful due to convergence," he says, "but Eldo predicted the conversion time exactly the first time." That was a

year ago; today Manning boasts, "We've never had a failure with Eldo due to convergence."


Jacques Benkoski, manager of design solutions development at SGS-Thomson Microelectronics (Grenoble, France), became a supporter of AnaCAD's Eldo about two years ago, after comparing it to his company's own Spice simulator and other Spice simulators. "It's the only simulator on the market that can cover what we're doing," says Benkoski. "We're a multinational company doing a large range of ICs—from purely digital to purely analog and everything



Using Eldo's XELGA graphical user interface, Dr. John M. Manning (left), president and CEO of Monolith Technologies, views with Akhtar Ali, a senior member of the technical staff, the results of simulating the A-D converter portion of their mixed-signal design. The multimedia chip being designed combines 20-bit analog performance with digital signal processing. Since it's aimed at the Japanese market, Monolith designers were up against a five-month design cycle. The chip is due to sample this month.

analog mode. Appropriate settings of the time-step and accuracy parameters for each algorithm can result in a significant overall performance increase. Benchmark tests in circuits ranging from 100 nodes to 50,000 nodes show an average improvement of 30 percent over previous versions of the Eldo simulator."

The enhancement to the FAS analog behavioral modeling module in Version 4.1.0 gives you the ability to quickly develop and verify complex analog behavioral models using the new FAS source-level debugger. Once the behavioral model has been veri-



The rapid proliferation of mixed-signal ASIC designs makes the availability of better EDA tools a major concern. Development of more effective tools for various

mixed-signal applications, such as "big D/little A" and "big A/little D," depends on the resolution of several issues.

These include hardware description languages (HDLs), synthesis, mixed-mode/level simulation, automatic place-and-route, and test.

■ Hardware description languages

The complexity and size of mixed-signal circuits make HDLS important for several reasons. Top-down design using high-level descriptions to drive synthesis tools will become the standard design practice. This methodology is widely used today for digital designs. Another important use for high-level descriptions will be behavioral simulation models in mixed-mode/level simulations. VHDL and Verilog are used extensively as digital HDLS.

Analog HDLS are another story. Characterization of analog circuits is quite different from digital. It's important that second-order non-idealities intrinsic to analog operation be included, along with the circuit's intended function.

A major problem here is lack of a recognized AHDL standard. Some commercial behavioral simulation languages are available—MAST, for example, from Analogy—and numerous companies have modified Spice for linkage with either C or FORTRAN models. It's currently unclear whether the AHDL standard, when it comes, will be an extension of MHD, an extension of VHDL or something else.

■ Synthesis

Logic and analog synthesis are a must for quick turnaround mixed-signal ASIC design. Logic synthesis and optimization tools are now widespread in digital design. Currently, logic synthesis maps a VHDL or Verilog description of an RTL description into logic gates. A future research area for logic synthesis is to extend the description level for synthesis from the RTL to the behavioral level.

Analog synthesis, on the other hand, is still mainly confined to university R&D. Most of the research is aimed at synthesizing fundamental analog blocks such as operational amplifiers, comparators, switched-capacitor filters, and D-A and A-D converters. Probably the most useful results have been in switched-capacitor filters, where it's relatively easy to put together predefined blocks. But the jury is still out on whether full automation is achievable for all analog circuits.

■ Mixed-mode/level simulation

Mixed-mode/level simulation capability is one of the more obvious needs of most mixed-signal ASIC designers. For clarity, mixed-mode simulation includes two different simulation domains: logic (gate-level) and circuit (transistor-level). Mixed-level simulation has a single simulation domain, such as logic, with both gate-level and behavioral-level models.

While many claim to have this capability, there often seems to be some simulation deficiency that frustrates the designer. What's needed is a robust native simulation environment that permits mixed simulation with transistor-, macro- and behavioral-level models, as well as charge-based, switched-capacitor elements having random inputs and clocks, z-domain blocks, s-domain blocks, and an event-driven engine for small logic blocks.

A key component for all areas of mixed-mode simulation is the interface models connecting different simulation domains—for example, logic and circuit. These are best developed as analog behavioral models that accurately characterize loading, input/output wave shapes and forward and backward time-step control.

■ Automatic place-and-route

Automatic place-and-route of standard digital blocks is probably the EDA capability that's made the digital ASIC business possible. The current major research area for digital is timing-driven place-and-route—that is, as feature sizes get smaller and clock speeds get higher, the digital layout problem becomes more analog in character.

The goal is to intelligently extract parameters so that timing information for critical nets can be either calculated or simulated to drive the place-and-route

system. Unfortunately, because of the significant amount of analog circuitry, the situation for mixed-signal place-and-route isn't as good. The typical digital place-and-route objective is to minimize chip area, total interconnect length (preferably along the critical path) and the number of vias. This isn't adequate for analog circuits because of critical constraints such as crosstalk, matching and thermal symmetry. It's clear that for successful mixed-signal ASIC design, advanced placement, routing and compaction tools that can account for critical layout parasitics, coupling and constraints are required.

■ Test

While digital test is difficult, analog test remains an art. There's been much work in the digital area, resulting in numerous test methodologies—for example, SCAN, BIST, BILBO—as well as ATPG software.

Analog circuits are usually tested by testing all of their specifications. This is becoming prohibitively expensive for mixed-signal. Future areas to address include better fault models and synthesis-for-test. Synthesis-for-test is key in that the designer will know that the finished circuit is testable.

It isn't unusual for the test engineer to have silicon in hand before starting to develop the test program. Clearly, a virtual test environment—Dantes from Cadence, for example—that promotes test development, debug and load board design concurrently with circuit design is required. Unfortunately, this does nothing to alleviate the need to test all specifications.

Probably the greatest obstacle to analog test is the lack of an adequate fault model. In particular, a parametric fault model is required, one that detects small variations from ideal behavior that cause the circuit to violate one of its specifications. A useful analog fault model must be developed that can help reduce both test development and test execution time.

W. Terry Coston, director of design systems, Harris Semiconductor, Melbourne, FL

Mixing signals and supplies

Systems and mixed-signal designers attempting low-voltage, low-power ASIC designs will soon find, if they haven't already, that there are some major design barriers their current development tools can't surmount. When mixing 3-V and 5-V circuitry and supplies, you'll come up against very distinct design software requirements. Here's a quick glance at some of the particular tool needs you'll face.

| | |
|--------------------------|--|
| Schematic capture | = cell-naming conventions to track voltage levels = netlist-auditing capability to ensure proper 3-V/5-V interfaces |
| Partitioning | = automatic reallocation of 3-V and 5-V sections for hierarchical circuit manipulation |
| Simulation | = software for optional V _{DD} assignment = full characterization of analog/digital cells at 3-V and 5-V = high-level analog block models for system simulation = single-pass timing calculation |
| Layout | = automatic routing of separate buses for 3-V and 5-V = floorplanning software to assign various sections to specific areas |
| Testability | = isolation of noise-sensitive analog blocks from digital = software to test for proper I/O voltage levels and dc characteristics = ATPG and production test for analog elements |

Contributed by **John Harrington**, department head, domestic OEM, ASIC design department, AT&T Bell Laboratories, Allentown, PA

in between. Eldo replaces Spice, covers behavioral simulation and handles 100,000-transistor circuits."

Eldo is also used by SGS-Thomson's analog semicustom division and at its design centers, where dozens of customers do their own mixed-signal designs. Benkoski also advises key accounts doing custom design at their own sites to use Eldo.

For digital parts of mixed-signal designs, until now SGS-Thomson has used its own Mozart simulator, with Mozart/Eldo for mixed mode. "We have moved to Verilog since it's an industry reference, and we'll use Verilog/Eldo for mixed-mode," adds Benkoski, who hopes in the future to propose VHDL for mixed-mode, too.

To check interfaces between analog and digital, SGS-Thomson uses Eldo/Eldo-XL for up to 100,000 transistors and Verilog/Eldo for larger circuits. "The biggest problem today is not simulation," claims Benkoski, "but the human interface for mixed-mode simulation. After designing a schematic for a circuit, how does a designer know what to simulate as analog and what to simulate as digital?"

SGS-Thomson is working with Cadence and AnaCAD on Verilog/Eldo to

ease that problem. Benkoski says, "It's up to the software to help the designer bring relevant information into the system." Among the issues involved are netlist splitting, translation of stimuli into analog and digital, single waveform display, and a


method of automatically inserting interfaces between the two simulators."

The absence of a standard AHDL also presents a problem. "Because there's no standard AHDL, there's no library of components. Nobody can agree on a language," says Benkoski. As for AnaCAD's FAS and Analog's MAST, he says they're very similar. "They're different in syntax and semantics, but globally they're the same. Whatever is possible with one is possible with the other, but it's not an easy problem to write analog behavioral models."

According to Benkoski, it's not possible to take a netlist, translate it and get a behavioral model. Nor does he agree with what he claims is Analog's proposal of extracting parameters from Spice to characterize a model. "If you don't write a model from the top down," he says, "it won't run any faster than a Spice model."

■ What about Profile?

But what about Cadence's Profile, which combines a behavioral expression language with a graphical environment to ease the task of model writing for system designers? Benkoski, a Cadence user, says he agrees that graphical views of behavioral components would be useful, since they could run on any simulator. But he wonders if Cadence will tie Profile, now linked to Valid's old Spice simulator, exclusively to Spectre. He asks, "Will Cadence



"Although analog tools have improved, they still aren't as sophisticated as digital tools," says Linda Eaton, a design engineer from GEC Plessey Semiconductors.

"Analog tools still require the engineer to analyze the circuit performance, where digital simulators can flag errors against expected results." She is making modifications to the layout for an advanced mixed-signal circuit for hard disk drives using Cadence's Analog Artist.

Until now,
there was just
one thing missing
from CPU32
development
solutions.



BUSINESS REPLY MAIL

FIRST-CLASS MAIL

PERMIT NO. 140

REDMOND, WA

POSTAGE WILL BE PAID BY ADDRESSEE




Applied Microsystems Corporation

5020 148th Ave. N.E.

PO BOX 97002

REDMOND WA 98073-9983

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES



The solution.


Introducing our newest development systems for the 68330, 68340 and 68F333 micro-controllers: emulators, CodeTAPs®, debuggers, simulators, language tools, and support services. Now you have a range of tools to match your needs throughout the development cycle. Only from Applied.

Smarter than your average tools.

Because these new 32-bit controllers are far more complex than their predecessors, you need a smarter solution. One that saves you weeks of development time.

For example, since you can reconfigure CPU pins, you have to be able to track the changes. Ours is the only system that can determine the correct function of each pin, transparently, and in real-time.

And our new intelligent trace disassembler provides the only accurate picture of register values correlated with instructions. It also gives you detailed timing



data for performance analysis.

Applied's system solves one of your most critical productivity issues. Instead of mismatched tools that don't speak the same language, you can now work with integrated tools that share a common user interface.

A real solution costs less than you think.

Just because our CPU32 development tools are the best in their class doesn't mean you have to pay more

for top performance. You'll be pleasantly surprised at how affordable these tools are.

Want to learn more?

We'll send you free product information and our latest Application Note on working with the complexities of the CPU32 controllers.

Just call Applied at 1-800-343-3659 (in WA, 206-882-2000) or return the card below.

And ask for your solution.

©1992 Applied Microsystems Corporation. All rights reserved. Other names indicated by TM and ® are registered trademarks of their respective holders. In the U.S. and Canada, call 1-800-343-3659 (in WA, call 206-882-2000). Or contact Applied Microsystems Corporation, P.O. Box 97002, Redmond, WA 98073-9702 USA. For the name of the nearest European distributor, call 44-296-625462. Fax 44-296-623460. Or contact Applied Microsystems Corporation, Ltd., AMC House, South Street, Wendover, Aylesbury, Bucks, HP22 6EF, England. In Germany, call 49-6182-92030. Fax 49-6182-920315. Or contact Applied Microsystems GmbH, Dammstrasse 6, W-6453 Seligenstadt, Germany. In Japan, call 03-3493-0770. Fax 03-3493-7270. Or contact Applied Microsystems Japan, Ltd., Nihon Seimei Nishi-Gotanda Bldg., 7-24-5 Nishi-Gotanda, Shinagawa-Ku, Tokyo 1141, Japan.

Send for FREE technical information.

Or call us at 1-800-343-3659 (In WA, 206-882-2000)

- ☐ Please send me product information about Applied's new development solutions for the following microcontrollers: ☐ 68330 ☐ 68340 ☐ 68F333
- ☐ Please send me your latest Application Note about the Motorola CPU32 microcontroller family.
- ☐ Please call me to arrange a demonstration.

Phone number _____ Title _____

Name _____

Company _____

Address _____ Mail Stop _____

City _____ State _____ Zip _____


CD1092



**Applied
Microsystems
Corporation**

Where it all comes together.

Both speed and accuracy in mixed-signal ASIC simulation? Absolutely!



Many mixed-signal ASIC designers are facing a dilemma. New manufacturing processes are letting them create ICs so large that simulation is becoming extremely

difficult. ASIC designers have an additional problem that manufacturers of standard ICs normally don't face—ASICs are typically designed concurrently with the target system. This often results in specifications changing during the ASIC design process.

Under these conditions, simulation needs to be more than just a verification tool. It must help the engineer iterate through multiple design ideas while developing a complete functional system. Without methodologies that drastically cut simulation times, mixed-signal design cycles can become painfully long.

Digital ASICs alone are pushing logic simulators to their limits. If analog is added to the ASIC, simulation times can increase to such a level that simulation may become virtually impossible. Normally, the analog portion of the circuit will ultimately determine how long it will take to simulate a mixed-signal design. A single Spice-level transistor, for example, can take more simulation resources than dozens of gates simulated digitally.

■ A better method of simulating

Because of these limitations, simulating the analog portion of a mixed-signal ASIC must be done with a different methodology. Partitioning the circuit into larger behavioral blocks, rather than into individual transistors, is an efficient method of greatly reducing simulation times with a minimal loss in precision. This method doesn't mean that ideal transfer functions or specialized controlled sources are used to imitate general circuit behavior. Instead, models are created that exhibit all the relevant behavior of a functional group, including secondary effects.

The next question is how best to divide the ASIC into logical functional groups. A reasonable delineation for these functional groups are the layout cell boundaries.

Creating these functional models can be extremely difficult if a traditional macromodel approach is used. Creating macromodels means combining a limited number of primitive elements in such a way that the equations that describe the behavior of the block are imitated.

A more direct approach is to use an analog hardware description language (AHDL). Essentially, an AHDL explicitly describes the electrical behavior of a device. Models created with an AHDL provide extremely fast simulations and yet still replicate detailed circuit behavior. This is because the designer can create a model that shows the desired behavior without any extraneous calculations that needlessly increase simulation time.

■ Mixed-signal model library

By using the capabilities of a powerful AHDL, an extensive library of behavioral models for ASIC cells can be created. These models can easily be parameterized to match a large spectrum of different cells, unique to each ASIC manufacturer's process.

Mixed-signal ASIC cell models can be categorized into four groups: digital, analog, mixed-signal, and analog/digital interfaces.

Digital models beyond the normal primitives may be modeled in a language available with the digital simulator portion of a mixed-signal simulation scheme. Such languages include VHDL and Verilog HDL. If the analog simulator has an AHDL that provides digital constructs, however, many of the simulations may be performed natively, where only one simulator is needed for both the analog and digital. Analog's Saber simulator with its MAST AHDL provides both native and mixed-simulator methods.

Analog modeling is needed for devices such as opamps, VCOS and phase comparators. Of course, all Spice-level primitive parts must be available so any section of the ASIC can be represented by individual transistors. It's important that even transistor models be represented by the AHDL and not be part of the simulator source code. Otherwise, it will be difficult to modify, or create additional transistor models as new technologies become available.

A large number of models needed for ASIC cells fall under the category of mixed-signal. This means that they have both analog and digital characteristics. Devices include A-D converters, D-A converters, comparators, and digitally controlled analog switches. Only an AHDL that recognizes both analog and digital capabilities lets these models be created.

An important class of models that's needed to accurately represent mixed-signal circuits is interface models between analog and digital. In Analog's Saber simulator, these models are called Hypermodel interfaces. These devices don't physically exist on an ASIC, but they're necessary to represent the electrical characteristics at the inputs and outputs of digital gates, so there's accurate interaction between the analog and digital models. Like other types of models, this interface may be parameterized to uniquely match each manufacturer's gate characteristics.

■ Expanding capabilities


By significantly reducing the simulation time, several new methodologies become useful for mixed-signal ASIC designs. With faster simulations, verifying the ASIC in the target application is possible. This capability can help eliminate the awkward situation where the ASIC meets specification, but it still doesn't operate properly in the target system. Of course, combined ASIC and system simulations such as these are only possible if the simulator has an abundant set of board-level components available.

Another capability of a mixed-signal simulator is the ability to concurrently simulate mechanical devices such as motors and transducers. Combining the ASIC, board and system into one simulation will greatly increase the likelihood of first-pass success.

Using an AHDL for mixed-signal ASIC cell models results in simulations that take minutes instead of hours or days. With a set of models already created, instituting such a methodology can be done fairly quickly. Faster simulations also make additional analyses, such as statistical, temperature and parameter sweeps and device faults, feasible. Using this methodology for mixed-signal ASIC design will result in better devices and shorter design cycles.

Mark Chadwick, BSEE, University of Wisconsin, Analog, Beaverton, OR

Place-and-route in mixed-signal ASIC and IC designs



cal design (place-and-route) of ICs. To increase system performance while lowering product cost, designers must put more logic on single, highly integrated chips.

To ensure correct timing in these circuits, designers have to influence the physical layout and gain greater control over the place-and-route portion of the design cycle. They can't afford to ignore physical design, passing a netlist over the wall to layout designers or the ASIC foundries; both lack full understanding of the original designer's intent.

■ Mixed-signal layout issues

Mixed-signal designers are concerned with the same performance and cost problems as their digital counterparts. In addition, mixed-signal designs present a number of specific issues in the domain of physical layout that designers must deal with. Among them are:

- Partitioning of analog and digital blocks. To limit signal interference or noise between analog and digital functions, designers must be able to easily segregate the analog and digital sections of the design.
- Multiple power and ground signals for analog and digital circuitry. When digital circuits switch, they cause numerous spikes in their power and ground wires. To avoid crosstalk with the sensitive analog circuitry, the digital and analog sections must have different power and ground signals. Mixed-signal chips typically use a large number of supply signals.
- Tapering of power supply signals. The width of supply signals should be small enough to guarantee minimal chip size, but wide enough to prevent metal migration and excessive voltage drop. This trade-off must be carefully balanced for the design to meet cost and performance requirements.
- Flexible wire width and routing rules.

Signals between analog blocks are usually wired in 1.5 layers (one metal routing with restricted polysilicon routing) or 2 layers, while the digital sections, because of their much higher clock speeds and the need for shorter delays, are wired in 2, 2.5 or 3 layers of interconnect with different line widths and spacing rules. Support for different wiring strategies and wire widths is critical.

To effectively address and manage these issues, mixed-signal designers require the following capabilities in their place-and-route systems:

- For minimal die size, optimal circuit performance and the highest degree of integration, mixed-signal place-and-route tools must be able to combine cells and blocks and analog and digital circuitry in a flexible floorplan.
- Although there are a few very common supply schemes (planar, non-planar, supply ring) that can be automated, not every possible supply routing strategy can be embedded in a place-and-route tool. For maximum flexibility and layout optimization, the tool must have facilities for interactive, symbolic supply wiring as well as the capability to complete supply signal routing automatically. With symbolic supply routing, designers can easily influence the layout without actually performing detailed manual layout tasks.
- The layout system must provide a means to analyze voltage drop and current density of any supply topology (including multiple supply pads), and then automatically taper the supplies to limit these drops and densities.
- Most analog blocks are designed with a full-custom approach, using total freedom in pin locations, pin width and block size. Digital cells and blocks are often generated with cell and block compilers. To import such blocks into a mixed-signal chip design without making pin modifications, the layout tool must be completely gridless.
- For optimal circuit density, the layout system must be capable of triple-level metal routing, with efficient use of over-the-cell routing.
- Automated place-and-route systems must be able to produce the smallest

possible chip while satisfying chip performance criteria. Timing-driven layout from system design specifications with accurate feedback on actual chip performance is a must.

Many of these mixed-signal design issues can only be handled by the original designers, who have expert knowledge of the design intent. This helps explain the current shift in design methodology that's bringing physical considerations forward in the cycle.

■ Positive impact

With the same designers involved in both the design and the physical implementation, it's possible to positively affect time-to-market. While the design clearly impacts the physical implementation, production schedules slip when design changes are required to achieve performance or density goals. When both design and physical implementation happen in parallel—one aspect of concurrent engineering—most of these issues can be addressed early on.

Because of the shift in design methodology that's occurring today, more emphasis is being placed on top-down floorplanning for physical design. Floorplanning from the outset allows early silicon real estate and chip performance planning. It also provides mixed-signal designers with early insight into the issues of supply planning and analysis and analog-digital circuit segregation. Floorplanners and place-and-route tools that use a common routing engine throughout the design cycle can work together to give designers better results with fewer design iterations.

To appeal to systems and mixed-signal designers, place-and-route tools must communicate in a language that's comfortable for engineers, rather than traditional layout specialists. Design specifications must drive the place-and-route process, and designers must influence the layout with the place-and-route tool, relying on their automatic algorithms to take care of the details.

To close the loop, designers need accurate feedback on how the place-and-route tool performed with respect to the original specifications. Only then can designers take advantage of a new design methodology that squarely addresses time-to-market.

Guido Arnout, vice-president of engineering, Silvar-Lisco, Sunnyvale, CA

REFLECTING ON THE POSSIBILITIES



*Alliant • Motorola Delta Series
• Data General • PC-ATbus • ISAbus
• EISAbus • IBM Micro Channel®
• DEC TURBOchannel® • Harris Night Hawk
• Concurrent • Silicon Graphics • Sun Sparc
• Encore • VMEbus*

VMIC'S REFLECTIVE MEMORY NETWORK BREAKTHROUGH

Very Fast Real-Time Communications Among the Widest Variety of Computers in the Industry. The Reflective Memory concept provides a very fast and efficient way of sharing data across distributed computer systems.

VMIC offers fiber optic and flat cable Reflective Memory Networks providing up to 256 nodes and up to 20 Mbytes per second transfer rate.

DEC and TURBOchannel are Registered Trademarks of Digital Equipment Corporation. Micro Channel is a Registered Trademark of IBM.

1-800-322-3616



**VME Microsystems
International Corporation**
12090 South Memorial Parkway
Huntsville, AL 35803-3308
(205)880-0444 FAX (205)882-0859

VMIC products are internationally represented by distributors throughout the world. Call or FAX VMIC for complete information.

spectre, or will they open it?

Dr. Gordon Cumming, vice-president and general manager of the digital communications and signal processing business unit at TriQuint Semiconductor (Beaverton, OR), says that TriQuint uses analog behavioral modeling for some applications. One issue for TriQuint is integrating TriQuint's enhanced gallium arsenide FET model (code name TOM, TriQuint's Own Model) into the analog simulation engines.

video sensor controller for an ATE project, senior staff engineer Jeffery Gutsell of Hughes Aircraft (Canoga Park, CA) had to develop his own analog and digital libraries for a new BiCMOS process from National Semiconductor (Santa Clara, CA) that he's using. The ASIC he's designing includes a 14-bit A-D converter, three 8-bit log current D-A converters, clock drivers and clock generators, a digital controller, and a high-speed parallel-to-serial and se-

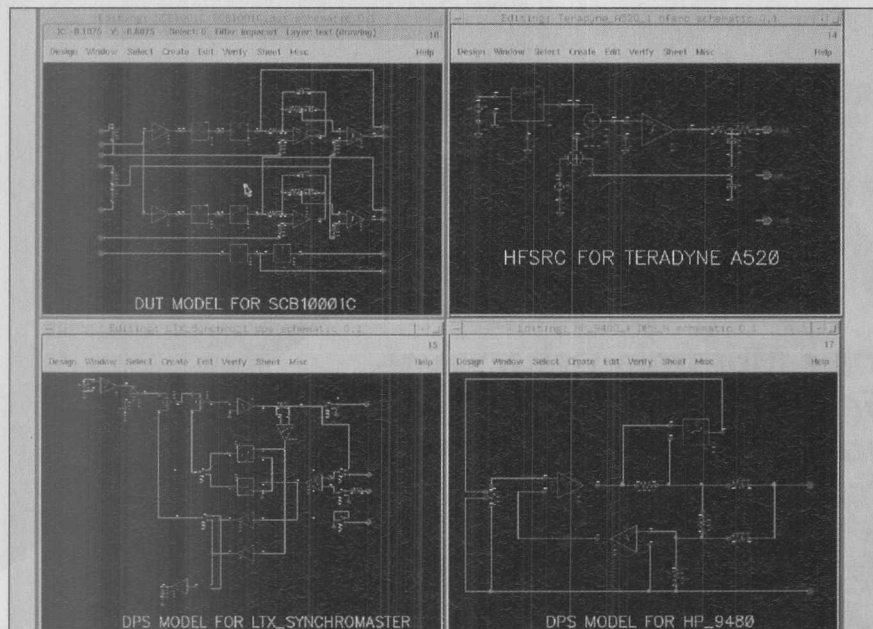
quential-to-parallel converter. The Design Architect schematic editor and the AccuSim analog circuit simulator through the Falcon Framework.

Mentor's mixed-signal design solution is based on its Explorer LSim multilevel simulator, which has a Simulation Manager backplane that lets you combine simulation algorithms from circuit to system level. Mentor recently made its new Scap switched-capacitor simulator available as an optional kernel to Explorer LSim. Because of LSim's multikernel simulation capability, Scap users don't have to debug digital and analog circuitry separately, nor must they check the interface manually, according to Dr. Edmund Cheng, director of strategic programs for Mentor's analog/mixed-signal division. "Many mistakes in prototype silicon are due to interface errors," says Cheng. "The solution is to debug the analog and digital circuits together."

More layout choices coming

Through its Analog Interface Kit, an interactive and graphical simulation environment, Mentor has recently integrated AnaCAD's Eldo circuit simulator into the Falcon Framework. "The integration of Eldo into Mentor Graphics' design environment lets our customers simulate very large circuits with superior convergence and accuracy," says Dick Akers, director of marketing for the company's analog/mixed-signal division.

Akers suggests that Mentor is now directing its attention toward a front-to-back-end mixed-signal toolset, a more complete offering than its IC Station. He adds, "Right now, the marketplace needs integrated mixed-signal tools. The current offerings are really stand-alone tools when you get a good look at them. Integration via a framework solution offers users flexibility in analog and mixed-signal design by giving them access to a complete set of tools, ranging from design creation through synthesis and simulation to IC layout. A successful framework should permit the coordination of layout tasks with front-end design tools, so users can visualize and match layout paths to electrical design nets. This allows back-annotation and simulation of the electrical database and increases confidence in the final design."



The first commercially available testability product for analog and mixed-signal design, Cadence's Dantes Design and Test Engineering System includes tools to simulate the device under test (DUT) and the tester environment, as well as providing tester resource modeling. The upper left window of this Dantes screen shows a model of a DUT. The remaining three windows show tester resource models provided by several ATE vendors. The upper right window shows the model for the high-frequency source for the Teradyne A520. The lower left window contains the DUT power-supply model for the LTX SynchroMaster. The lower right window shows the DUT power-supply model for the HP 9480. Dantes also includes design-for-testability tools, tester rule checking and ATPG.


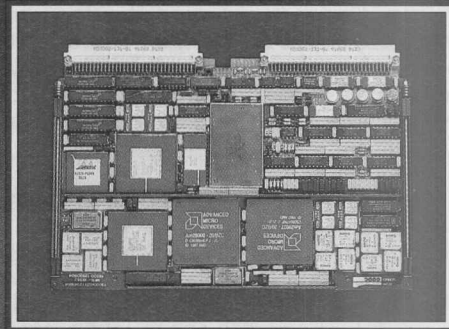
"There are GaAs FET models available," says TriQuint's CAD manager, Marty Zimmer. "However, TOM is an enhanced model." The last time Zimmer evaluated mixed-signal environments, the tools that were best from a mixed-signal perspective were among the worst from a model-refinement perspective. Since then, however, several of the companies specializing in analog tools have incorporated refinements for TriQuint's TOM model. "Major vendors like Mentor and Cadence are slower to pick up models like these than vendors which specialize in analog simulation," says Zimmer, "but they're starting to respond."

rial-to-parallel converter with ECL interface that runs up to 300 MHz.

A Verilog/Saber user, Gutsell used MAST to write analog models at the behavioral level. Anything that was critical he put into Spice. "Saber is very flexible," he says. "You can make a model as sloppy or as accurate as you want. The biggest problem is that Saber is very slow compared to Verilog. You wind up sitting there waiting for the analog forever."

It can also be a real chore to generate a netlist for simulation if you're going from one platform to another, according to Gutsell, who works on a Mentor platform. For

We're Proud to say we're Under the Gun.



The Fire Control Processor for the M1A2 Abrams Main Battle Tank: one example of Vista Controls' fully mil-qualified RISC-based board-level systems solutions for fire control and gun stabilization. Being manufactured today.

A Cost-Effective Solution...Now

If your system requires NDI (non-developmental item) MIL-SPEC VME/SAVA/1553 boards and chassis, then Vista Controls can supply your needs. Our RISC-based products typically take the place of two to three CISC-based boards, resulting in significant system cost savings. Vista's NDI hardware and software shorten your leadtime, lower your costs, and reduce your risk--let our RISC lower your risk.[™]

Validated ADA and C
SAVA and 1553
Radiation Hardness
Control Systems

Real-time Kernel
IFTE Test Interface
ATR Chassis
Real-time Communications


With Quality and Support Built-in

Vista Controls' engineering and software expertise in fire control, real-time communications, and control systems will provide you with complete turn-key systems solutions all the way from concept and systems integration to full production.


Whether your system operates on land, in the air, or on the ocean, Vista Controls stands ready to supply you with the highest performance VME cards available for fire control, communication, and control systems.

Ready to go under the gun...for you.


VISTA CONTROLS[®] CORP.
HIGH PERFORMANCE MILITARY CONTROL SYSTEMS
Valencia, CA 91355 • (805) 257-4430
Total Systems Integration[™]



NDI: Control and I/O Systems



NDI: FDDI and Ethernet Networks



NDI: Fire Control Systems/Auto-Loader



NDI: Gun Control/Auto-Loader

SPECIAL REPORT: ANALOG DESIGN

Mentor's forthcoming release of a more comprehensive mixed-signal offering—with layout tightly tied to other tools—will expand layout choices for those who've counted on Cadence until now for back-end capability. Most of Cadence's analog and mixed-signal customers are still using Analog Artist 2.4, based on the original Cadence framework, Design Framework I—also referred to as the Edge product line.

When Cadence released Analog Artist 4.2, based on Design Framework II, last January, it didn't include mixed-signal design capabilities. If you're a Cadence user

Falcon Framework, Viewlogic's Powerview and the Dazix IC design environment. Compass Design Automation also offers layout tools for mixed-signal design.

Still a lot of handwork

Hughes Aircraft designers are currently using Cadence's Edge product line to complete the not-very-automated task of laying out their chip. The group developed some of its own tools just to ease the analog burden. It did the analog portion by hand with Cadence tools on a Mentor platform, and the digital part was done with the Cadence Cell 3 Ensemble

where designers also use Cadence's Edge products at the back end and Viewlogic's schematic capture and ViewSim with Meta-Software's HSpice at the front end. "With digital layout, we rely on tools to do 90 percent of the work, whereas with analog, the tools do 50 to 60 percent if we're lucky," Nekar says.

Expectations and reality are two different stories with tools, claims Nekar, "especially when interfacing one set of tools for front end and another for back end. It took six months to get the tools operational and to get a die size that was palatable." Exar now has Artist Overlay, an enhancement from Cadence that lets you do mixed-mode layout more easily.

Mixed-signal mostly turnkey

Except for its strategic accounts, Exar discourages users from designing their own mixed-signal chips at design centers or at their own sites. "To technically support an analog design environment is probably the biggest nightmare," says Nekar. "No matter how general-purpose you've made an analog building block, you're going to have to tweak it. System designers won't have a good idea of what the silicon's going to do. We could be misleading them to believe that what they're designing is going to work in their application the first time."

Although GEC Plessey Semiconductors (Scotts Valley, CA) has made Cadence's Analog Artist 2.4 toolset available to customers with libraries, only a few users have purchased the tools, according to Brent Wilkins, marketing manager of military, industrial, and automotive products. It's been over a year since Plessey switched to Analog Artist internally, and there's been a very long learning curve, but after using a variety of tools from different vendors, Wilkins reports that engineers are very pleased with a single database on a single platform instead of having design data spread across different systems, requiring translations and paving the way for the interjection of human error.

Workstation CAE supervisor Ed Gilbert, who's responsible for bringing up Plessey's bipolar libraries on Cadence's Analog Artist, says that engineers are benefiting from the ability to transfer the database from one phase of design to another, and also from the speed with which this can be done.

AHDL standards activities update

IEEE-1076.1 covers the analog extensions to VHDL. Although VHDL is scheduled to be a standard language in 1994, two VHDLs will actually exist: one with analog extensions and one without. The analog extensions will be merged back into IEEE-1076 in 1997. Although IEEE-1076 and 1076.1 are entirely volunteer efforts, the U.S. Air Force's Rome Labs has recently put out a request for proposals that will support the analog extensions to VHDL and fund a prototype.

On the MHDL (MIMIC hardware description language) front, the DOD has contracted with Intermetrics (Cambridge, MA) to design an HDL for the MIMIC (microwave monolithic integrated circuit) program. A first draft of the language reference manual (LRM) has been released, and sample models are being written. The final LRM will be completed in early 1993, and a prototype will be developed after the LRM is released. Once the feasibility of the language and its implementation are established, the LRM goes to the IEEE for standardization.

Contributed by **Mark Chadwick**, BSEE, University of Wisconsin, Analogy, Beaverton, OR

designing mixed-signal chips, you have to wait until this December to upgrade your design environment. At that time, Cadence will make Analog Artist 4.2.1A available; it's built around Design Framework II, with mixed-signal layout capabilities, including the Virtuoso digital layout tool with analog extensions and Cadence's new Preview floorplanning tool. The new release will also include improved Spectre integration and Spectre/Verilog-XL.

Cadence users who have the Design Planner II floorplanning tool, developed by High Level Design Systems (Santa Clara, CA), have another alternative—to use the SC place-and-route engine from Silvar-Lisco (Sunnyvale, CA) for designs combining analog and digital blocks and cells. SC has been directly integrated into the database of Cadence's Design Frameworks I and II under the Design Planner II floorplanner. Users can also access Silvar-Lisco tools through Mentor's

autorouter on Sun. It was up to the designers themselves to integrate the analog and digital, according to Hughes' Gutsell.

Monolith's Manning chose to use Compass tools for layout, not because they're the best, he says, but because they're very complete and work well with AnaCAD tools. "Compass layout tools have gone from cumbersome to still cumbersome," Manning adds, "but the toolset is so complete, going from silicon compilation to layout and extraction. The tools play well together until we start pushing them with our mixed-signal designs."

Although Monolith designers have a Compass floorplanner tool, they didn't use it. "We did it all by hand," says Manning. "There's still a lot of handwork for high-precision analog."

Layout is the major bottleneck of mixed-signal design, according to Nejo Nekar, director of strategic products at Exar (San Jose, CA),

16 lines, No waiting.

Your application is first in line with the MVC 16-line Async Commux. It's got processing capacity to spare, thanks to a 16 MIPS RISC, so system power goes to your users — not I/O.

Character processing and buffering is managed on-board by our RISC, so driver calls and host overhead are kept to a minimum. All 16 lines can operate at 38.4 KBaud. That's over 61,000 characters per second throughput, double the rate of other VMEbus async controllers.

The MVC's advanced features benefit both the integrator and programmer. Port and VMEbus parameters are soft-configured and set line-by-line. Modem control is standard. Full software support is also included,

along with diagnostics and a Streams driver.

Advanced memory architecture and high-speed buffering eliminate overruns and port domination. Memory is expandable from 128KB to 1MB, so the MVC handles today's requirements and future application needs.

With over 11 years experience producing advanced storage, communications and memory products, Macrolink delivers the powerful and flexible MVC in 8 and 16 line configurations. Call us today. We won't keep you waiting.

Macrolink Inc.,
1500 North Kellogg Drive,
Anaheim, California 92807.
Phone (714) 777-8800,
FAX (714) 777-8807.


macrolink inc.

CIRCLE NO. 64

INSTANT DATA ACCESS (IDA)

DIAL (617) 494-8338 DOCUMENT NO. 1037

*Mix & match I/O
panels feature
DB25 or RJ-45
connectors*



and Spectre in-house. For mixed-mode, designers use Cadence Spice with Verilog. "Since using Analog Artist tools, we've sent several chips to production," says Gilbert, "all 100 percent on first silicon."

■ Design-for-test software

Plessey doesn't yet have Cadence's

Dantes Design and Test Engineering System, expected to ship this month. The sudden enthusiasm for testability in the digital arena and the explosive market for test tools haven't reached the world of analog and mixed-signal, where Dantes is the first commercially available design-for-test product.

request of Analog Alliance partner Harris and worked with ATE vendors to implement it. At this point, Harris is using Dantes in-house but isn't offering it to users because it's brand-new. Dantes will be integrated into Analog Artist 4.2.1A, due from Cadence two months from now.

Dantes presents a graphical and automated approach to test program development, letting you begin test development early, by using simulation instead of waiting for silicon. With a common database for design and test, Dantes includes tools for design-for-testability, tester rule checking, tester resource modeling, simulation of the device-under-test and tester environment, and automatic test program generation. Tester models that interface with Dantes currently have to be purchased from ATE vendors.

SGS-Thomson is evaluating Dantes but, in the meantime, is using a homegrown program; the latter is similar to and based on Cadence technology, reports R&D manager Christian Caillon at SGS-

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

AnaCAD
(510) 656-0979 Circle 230
Analogy
(503) 626-9700 Circle 231
AT&T Microelectronics
(800) 372-2447 Circle 232
Burr-Brown
(602) 746-1111 Circle 233
Cadence Design Systems
(408) 943-1234 Circle 234
Comdisco Systems
(415) 574-5800 Circle 235
Compass Design Automation
(408) 434-7820 Circle 236
DAZIX
(303) 581-2300 Circle 237
Exar
(408) 434-6400, ext. 3586 Circle 238
GEC Plessey Semiconductors
(408) 438-2900 Circle 239

Harris Semiconductor
(407) 724-7000 Circle 240
High Level Design Systems
(408) 982-2550 Circle 241
Mentor Graphics
(503) 645-1551 Circle 242
Meta-Software
(408) 371-5100 Circle 243
National Semiconductor
(408) 721-5000 Circle 244
SGS-Thomson Microelectronics
(408) 452-8585 Circle 245
Silvar-Lisco
(408) 991-6100 Circle 246
TriQuint Semiconductor
(503) 644-3535 Circle 247
Viewlogic
(508) 480-0881 Circle 248

Get down without getting into

Look at the 3V experience of AT&T Bell Laboratories. Look at the success of our sub-3V cordless phone. You can see why no one can match AT&T's problem-free, worry-free 3V system design development.

The chart tells you why we're a 3-volt leader. But, beyond that, you get the design-in power of our flex-

| 3-Volt ASIC Checklist | AT&T |
|--|------|
| 1. Experienced 3V ASIC Supplier | ✓ |
| 2. > 70 MHz 3V Operation Achievable | ✓ |
| 3. Fully Characterized Library of Cells (Voltage/Temp./Process) | ✓ |
| 4. Mixed 3V/5V Chip Design, Audits, Layout | ✓ |
| 5. 3V Industry-Standard Macrocells | ✓ |
| 6. 3V Analog Cells (OSC., A/D, D/A, etc.) | ✓ |

the program sets out to improve the transfer of test information between design and test engineers, which is very important in semicustom, says Caillon, since designers are based all over the world. In the mixed-signal semicustom area, Caillon reports that 50 to 100 different circuits are generated every year.

In the TDS semigraphic software program, schematic capture of all test configurations, which uses simplified tester resources, is based on Cadence technology. Basic schematic capture, begun by the designer, is completed by the test engineer, who adds tester-related information. After a netlist is issued from the schematic, a translation program automatically generates a test program and a test board schematic.

TDS supports a large tester library, including both software and hardware data. This lets it translate all the captured information into a source test program. ■

It's a mistake to consider the progress made in analog design relative to state-of-the-art digital. The chasm between the two may never be bridged. But if analog design today is compared to analog design a year or so ago, a significant improvement can be seen—in the direction of the technology, if not in the technology itself.

The efforts going into the development of a standard AHDL give proof positive that the industry recognizes the significance of such a standard. Though not unified at present, those efforts deserve your support. And since competition is always a good motivator, the U.S. presence of AnACAD (whose FAS analog behavioral modeling option is an alternative to Analog's MAST) offers further assurance that technology in this area will not stand still.

The competitive environment developing around mixed-signal layout, up until now almost exclusively

out tools, now that you have a choice, consider the link from layout to other tools (including floorplanning) a priority, since that link's impact will loom larger and larger as time goes on.

As for analog testability, it may be creeping and crawling along, but even a single commercially available product is a start. After all, with analog design, we have to savor the joy that comes with every little growth spurt, knowing that somehow silicon will happen anyway.

Barbara Tuck Egan



to a cool 3 volts hot water.

That's AT&T "Customerizing."

ible, fully integrated ASIC libraries.

You can, for example, mix and match cells in our 3V library with cells in our 5V library. Or even choose CMOS or BiCMOS.

Clearly, AT&T is redefining ASIC by giving you the tools for success. Instead of the traditional "Gates by the Pound" ASICs supplier, we offer

you a silicon partner who understands your system. Who provides the leading technology and integration functionality needed for an optimum ASIC solution.

We provide a powerful boost to system design productivity with the absolute most in ASIC Macrocells, targeted at notebook PC, worksta-

tion, mass storage, cellular and multimedia applications. All characterized to run at 2.7V or 5.5V. All compatible with our three standard cell libraries in CMOS and BiCMOS.

For more about the support we offer in 3-volt development, just give AT&T Microelectronics a call at 1 800 372-2447, ext. 907.



AT&T
Microelectronics

We Define Performance

**INDUSTRIAL
SINGLE BOARD
COMPUTERS**

E486

**Full 80486DX 50 MHz
EISA Performance with
PCXI Compatibility.**

- 256K Write-Back Cache
- 128 MB RAM Capacity
- Drives 6 EISA Master, 8 EISA Slave Slots and 6 ISA Slots
- Keyboard & Mouseport
- Watchdog Timer • Dallas®

Real-Time Clock with 10 year battery

- 486DX 50 MHz and 486DX2 66 MHz now available

L486

**Full 80486DX 50 MHz
PC/AT Bus Performance
with 32-bit Local Bus.**

- Adaptec® SCSI-2 Compatible Interface • 2 Serial
- 128K Write-Back Cache
- 1 Parallel • FDC • IDE • Keyboard & Mouseport
- Watchdog Timer • Optional 512K FLASH • 32-bit local bus connector for feature daughterboards like SVGA, LCD, Solid State Drives, etc. • 486DX 50 MHz and 486DX2 66 MHz now available

Toronto MicroElectronics Inc.

- Industrial Single Board Computer Solutions for OEM's, VAR's, SI's and Embedded Applications
- Products Designed and Manufactured by TME exclusively in North America
- All products covered by full two year warranty
- Other available products include: 386SX 16/20/25, 386DX 25/33/40 and 486DX 33/50 w/o SCSI or Local Bus

ALL TRADEMARKS ARE THE PROPERTY
OF THEIR RESPECTIVE OWNERS.
SPECIFICATIONS SUBJECT TO CHANGE.
CIRCLE NO. 68

TME

TORONTO MICROELECTRONICS INC.

1815 MEYERSIDE DRIVE, UNIT 12, MISSISSAUGA, ONT., CANADA L5T 1G3
TEL: (416) 564-4833 FAX (416) 564-4768

Xerox brings cost-efficient strategy to The Reading Edge



The Reading Edge design team from left to right: Phil Minasian, manager, adaptive technology engineering; Richard Wheatley, diagnostics engineer; Matthew Schreiner, software engineering manager; Hunter Kennedy, senior electrical engineer; David Van Iderstine, senior systems software engineer; Larry Perletz, senior software engineer; Joanne Vining, senior engineering technician. Team member not photographed: Steve Baum, chief scientist. Minasian is holding the PCB containing all the electronics in The Reading Edge. The black box on the table is The Reading Edge machine itself. Its chassis has an angular shape to hold an open book or magazine for scanning.

A few years ago, a blind person wouldn't have been able to read this article or any other printed material. Then, in 1975, designers at Xerox Imaging Systems (XIS—Peabody, MA) invented a machine that could recognize characters on a printed page and convert them into synthesized speech. Since then, engineers at XIS have developed six generations of its reading machine, each time making it smaller, faster and more accurate, and adding more features while lowering the price.

At the heart of the device is software that performs intelligent character recognition (ICR), a technology upon which Xerox Imaging Systems was founded. The ICR algorithms recognize any print font, learn different attributes of a character, convert characters into words, and read the words aloud. The impact of this technology on the blind and visually impaired is tremendous. It increases their employment and educational opportunities, and improves their quality of life by enabling them to read printed material.

Early versions of the device were the size of a washing machine, weighed over 100 pounds and could be afforded for personal use by only the very few. "In response to customer demands," explains Greg Guidice, director of marketing for the product, "we

Jeffrey Child, Senior Editor

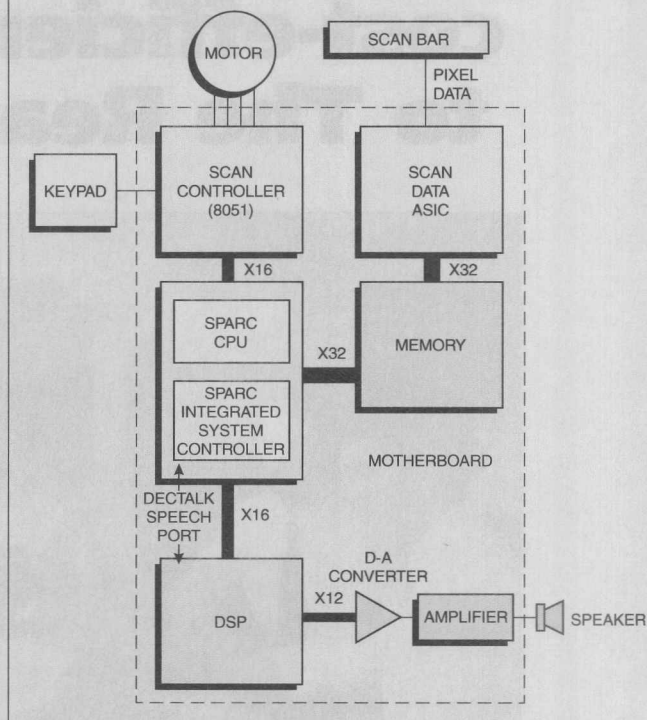
Designed for blind or visually impaired users, The Reading Edge is a machine that scans printed documents—books, magazines or any printed page—and converts the words into audio output. You lay a piece of paper or a page from a book on the machine, lining it up properly along the edge. The orientation of the page doesn't matter. You then press a key and the scanner head moves across the page.

As data is acquired and fed into memory, recognition software running in the SPARC turns the data into letters, which are then passed off to a speech task which sends phoneme parameters to the DSP. The SPARC processor turns the letters into speech, which then is passed off to the DSP and finally to the speaker.

From there, most of the processing is done in the software realm. A DSP is used to implement a set of filter functions. There's a noise source built into the DSP, and the filters follow it. The SPARC CPU sends parameters, called frames, that control how the filters behave. Each frame represents 6.4 ms of speech. Every 6.4 ms the filters in the DSP are modified slightly, resulting in a slight change to the sound that comes out the speaker.

Within about 15 seconds of aligning the first page, the machine will start reading it aloud. As soon as the machine starts talking, you can take the page off and put another one in its place. The second page is scanned while the first one is being read.

The Reading Edge electronics



wanted to create a portable, light-weight reading machine that you could use at home, at work, in the office, or take to the library. And because unemployment is high among the blind, leaving little disposable income, reducing the price of the machine was a key design goal." After a year of design work in which the Xerox team strove to meet size and cost goals, the company will unveil this month its most recent version, a machine called The Reading Edge. It costs under \$6,000.

■ All in one box

While The Reading Edge's immediate predecessor consisted of two separate units, a processor and a separate scanner, XIS wanted the new machine to be combined into one box. Matthew Schreiner, software manager at XIS, was responsible for the software design of the last two generations of the reading machine. "We knew," he says, "even as we were designing the previous model, that it was too complicated. There

were too many boards. It was too difficult to manufacture, with too many precision components that had to fit just right. We wanted to pull together the components and simplify them, so that we could implement the design on a single board. There were three specific guiding goals: one board, simple mechanics and a scanner head that moves back and forth."

Once the team decided to make an integrated box, the first step was to reduce the number of processors used in the old design. The previous generation had a 68020 as the main processor, an 8048 keyboard controller in the main processor unit and an 8048 in the keyboard itself, and an 80186 to drive a DSP chip for speech output. In addition, the separate scanner unit had its own processor. "There was a proliferation of processors," says Phil Minasian, program manager of The Reading Edge project. "Along with the idea of integrating the boxes, we wanted to do the same thing with the electronics. In

the new design, we have a single main processor. We still have a DSP, but now it's driven by the main processor instead of the 80186."

Not everything is controlled by the main processor, however. A separate 8051 8-bit microcontroller drives both the scanner and the keypad. "We retained a processor in that area," Minasian adds. "It was really a cost/performance trade-off. The 8051s are dirt cheap."

■ Choosing a chip

As the design team sought to reduce the number of processors, it faced its most critical decision in the choice of a main processor. The team had experience with both the Motorola 68020 and the Advanced Micro Devices 29000. As a result, it considered using AMD's 29005 RISC embedded processor, as well as embedded versions of the Motorola 68030 and 68040. But previous experience with a processor, although significant, was not as important to the XIS team as cost. "We wanted a cheap proces-

essor that had companion chips available, to minimize the amount of work we would have to do within the CPU architecture, and focus on the outside instead. We actually came very close to using the Acorn [a RISC CPU built by VLSI Technology]. It really fit from many points of view, especially cost."

tion, and hardware breakpoint capabilities. Unfortunately for XIS, along the way LSI Logic canceled the companion chip as a standard product. LSI Logic gave XIS the rights to the chip, but the company had to finish simulating it itself. XIS implemented the chip as a gate array.

"Overall, the companion chip has

custom chip would be needed. Hunter Kennedy, a senior electrical engineer at XIS, designed a data acquisition chip, dubbed the SCAN ASIC, in a 9,000-gate Xilinx FPGA.

■ Hard-working ASIC

Although Xerox already had tools from competing vendor Actel, and had used Actel ASICs in a previous product, Xilinx was chosen for cost and packaging. "It was \$400 for an 8,000-gate device from Actel and you couldn't reuse it," says Kennedy. "Also, Actel didn't provide an 84-pin package at that time."

The SCAN ASIC performs several functions in the system. The SPARC CPU sends a command to acquire a number of scan lines and informs the 8051 of this. The 8051 and a PAL which performs synchronization and CCD control send VALID line indications to this SCAN ASIC. After collecting 32 bits of 1-bit video data, the ASIC flags the main SPARC CPU, instructing it to pick up the collected data.

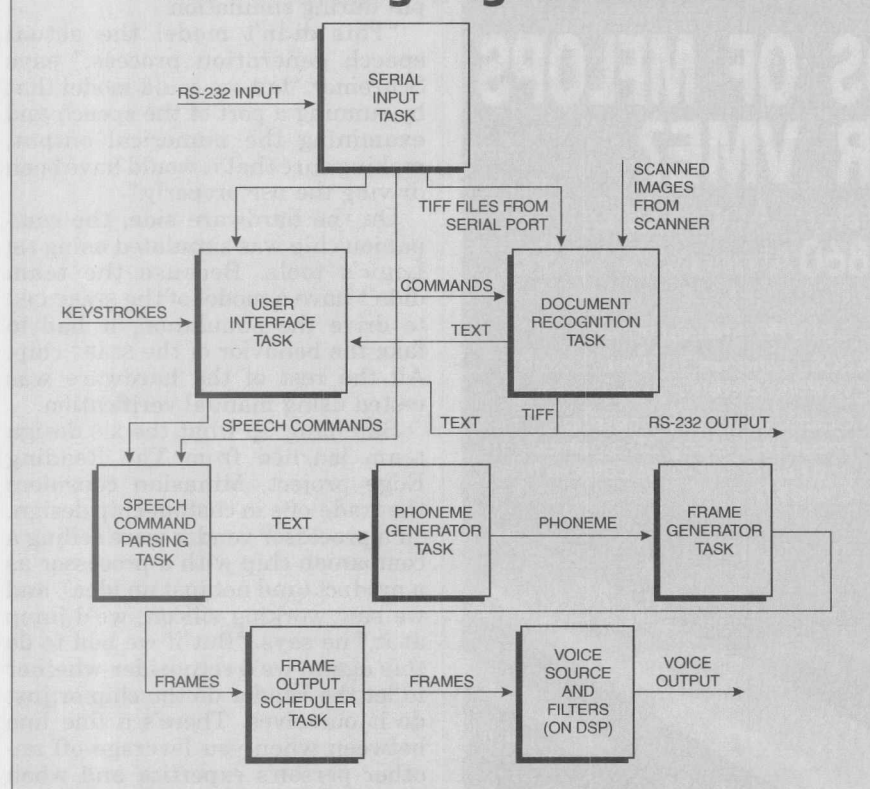
And data acquisition isn't the only job performed by the ASIC. "Since our system uses only black and white, we need to adjust the level of threshold to recognize black or white at the right position," says Schreiner. "Our SCAN ASIC drives the SRAM, which outputs adjusted thresholds to a DAC in the threshold circuitry. It also lets us calibrate for any non-uniformities in the response of the CCD's [charge-coupled device] pixels or the lamp [in the scanner]."

Because each pixel in a CCD has a different gain and bias, it was necessary to provide a threshold value for each pixel. With 2,600 pixels, this meant storing 2,600 threshold values. These values are stored in dedicated SRAM hanging directly off the SCAN ASIC. The SRAM is eight bits deep to permit dithering.

"This turned into a partitioning issue," says Schreiner. "For example, we could have put that SRAM inside the SCAN ASIC, but that would have been expensive. It made more sense to put it outside as a standard part and just have the ASIC deliver the addresses to it."

If the team had chosen to store the thresholds in main memory, it would have required another data channel to read them out into the SCAN ASIC. That would have meant another DMA channel and more load on the bus. So the decision was made to bring

The Reading Edge software



Much of the processing in The Reading Edge machine is done in software. Images are received from the scanner, converted to text and broken down into phonemes. These phonemes are divided into discrete time segments called frames. The frames are sent to the DSP, where they modify filters that control the speaker output.

In the end, the team chose the 901 embedded SPARC CPU from LSI Logic (Milpitas, CA) to be the main processor in the system. At the time, XIS was among a handful of companies that were helping LSI Logic define a companion chip for its SPARC CPU. Largely because of XIS's input, the companion chip was almost perfect for the reading machine design. It had all the necessary glue logic and peripheral functions the company needed.

Specifically, the companion chip provided integrated interrupt control, DRAM control, RAM refresh, two channels of DMA, bus control, cache

been successful, but it was a painful situation," says Minasian. "At the outset we had looked at doing the companion chip ourselves as an ASIC, but naturally we wanted to capitalize on the volumes associated with an off-the-shelf product. Cost was the key. At \$90, the cost of a two-chip set was far below anything we could compare it to. Again, while we had a lot of experience with the 68000 and 29000 architectures, there weren't companion glue chips available for them."

In keeping with its cost-conscious strategy, XIS used standard ICs wherever possible. But the special kind of

DESIGN STRATEGIES: IMAGING SYSTEMS

the SRAM outside. It was then a matter of partitioning the hardware so it was efficient, yet it had to make use of standard parts.

Ad hoc simulation

Software manager Schreiner stresses the role simulation played in enabling software development to keep up with the hardware. It was only in the last two months of the cycle that the hardware was actu-

ally working in its final form. Modifications were made up until the last week. Schreiner says, "We modeled as much as possible so that the day the hardware showed, we'd be able to implement the design."

For cost reasons, the team decided that full CAD simulation didn't make sense. "We had fully planned on using CAD and simulating the entire thing," says Minasian. "But we went straight to PCB. We

couldn't justify the cost."

Lacking CAD simulation tools, the XIS team adopted an ad hoc simulation strategy, simulating each component as best it could. The SPARC code was all simulated on a workstation. The code running on the 8051 was developed using an emulator. The one thing the team couldn't model was speech. To get around the problem, it attached the previous reading machine model via a serial port and sent text to it. In this way team members could at least listen to speaker output during simulation.

"This didn't model the actual speech generation process," says Schreiner, "but we could model that by running a port of the speech and examining the numerical output, making sure that it would have been driving the DSP properly."

On the hardware side, the companion chip was simulated using LSI Logic's tools. Because the team didn't have a model of the SPARC CPU to drive the simulation, it had to fake the behavior of the SPARC chip. All the rest of the hardware was tested using manual verification.

Summing up what the XIS design team learned from The Reading Edge project, Minasian considers the trade-offs in custom-chip design. "If a processor vendor were selling a companion chip with a processor as a product (and not just an idea), and we saw working silicon, we'd jump at it," he says. "But if we had to do this again, we'd reconsider whether to let the vendor do the chip or just do it ourselves. There's a fine line between when you leverage off another person's expertise and when you need to keep what's critical to your program in-house."

To keep cost under control, the XIS team decided it couldn't justify purchasing expensive CAD simulation tools. But what about next time? "Were I to do this project again, would I insist on CAD?" asks Minasian. "I think so. On the other hand, having been through this program, I've also seen the economy of what we've done. The only way we could even think about doing what we've done is because of the SPARC companion chip. Had that not existed, and had there been extensive glue circuitry required, then we absolutely would have needed simulation. But there's something really neat about a chip that integrates all the control functions. So if that chip works, and you've simulated that one chip, there's not a lot of hand verification left to do."

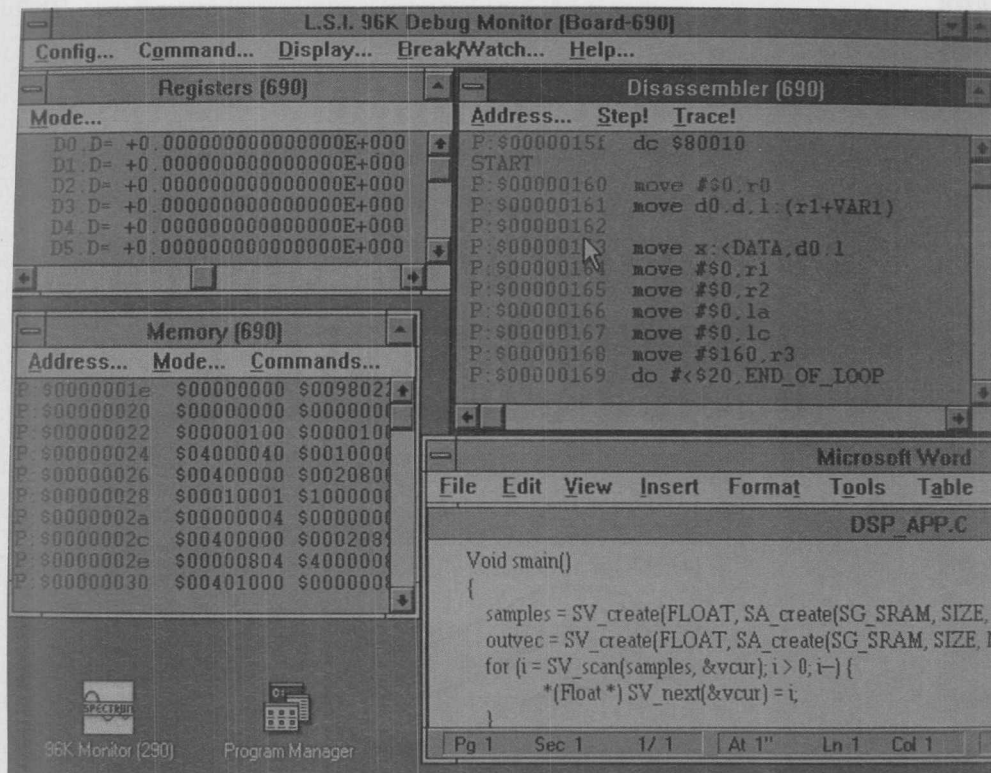
NEED HUNDREDS OF MFLOPS ON AT OR VME?

Put your Compute Intensive Applications on ALACRON's **Fastrack™ FT860** AT and VME Boards.

- ALACRON's open Fastrack™ inter-processor and I/O bus (200 MB/sec and 400 MB/sec options)
- Up to four i860XP 50 MHz processors per slot (200 Mips/400 Mflops per slot)
- 1K Complex FFT in under 300 microseconds (dual processors)
- Flexible high speed digital I/O options for: DT-Connect, ITI VISIONbus, SCSI, and Frame Buffer
- ALIX — ALACRON's multitasking, multiprocessing real-time OS
- C and F77 ANSI Compilers
- DOS / Windows 3.1 / Unix / Solaris / VxWorks Drivers
- Optimized signal processing, image processing, 3D graphics and numeric libraries callable from C and F77

CIRCLE NO. 69

Key commands accessible via simple pull down menus will give you the best view of DSP development.



You can view registers, data memory and disassembled program memory simultaneously.

When using more than one DSP board, multiple instances of Views™ may be invoked.

Edit and view Source Code through programs such as Microsoft Word for Windows.

If you think DSP is hard to develop, maybe you should go Windows shopping.

When you shop around, you'll find only one DSP solution that includes a Microsoft Windows 3.0 debugger interface.

The one from Spectrum Signal Processing.

Spectrum is putting DSP to work making application development quicker and easier than you ever thought possible.

By taking DSP software development to the next level of productivity, Spectrum now gives you the simplest, most complete solution to your DSP applications.

Our View96™ and View5x™ development solutions include everything from Windows monitor debuggers, to a host of software tools including C compilers, assembler/linkers, libraries and operating systems.

Plus you'll find bus and I/O flexibility, on-board A to D and D to A conversion, our DSP-Link™ for peripherals and much more. All backed by the complete technical support you've come to expect from the leader in DSP technology.

So if you've been skeptical about putting DSP to work, call Spectrum now and find out why everyone is looking into our Windows: 1-800-663-8986 (Western U.S.), 1-800-323-1842 (Eastern U.S.) or 604-421-5422 (Canada).



Our software selection includes all the tools needed to get you up and running fast.



Putting DSP to work

© 1992. Spectrum, DSP-Link, View96 and View5x are trademarks of Spectrum Signal Processing, Inc. All other trademarks are trademarks of their respective holders.

See us at DSPx booth #511

CIRCLE NO. 70

2ND HNNUHL ANALOG & MIXED-SIGNAL DESIGN CONFERENCE

ANALOG & MIXED-SIGNAL

10101001010010100100010011
00011010010000101011101001
10001001010101000100101010
01000101010010001010100111
01000101010000101010010101
11001100010100011011000010
01100010100110010110001101

DESIGN CONFERENCE

October 28 - 30, 1992

Hyatt Regency

San Francisco Airport

Burlingame, California



MAIL OR FAX TO:

Analog & Mixed-Signal
Design Conference
P.O. Box 7843
San Francisco, CA 94120-7843
Phone: (415) 905-2354
Fax: (415) 905-2220

ANALOG & MIXED-SIGNAL

10101001010010100100010011
00011010010000101011101001
10001001010101000100101010
01000101010010001010100111
01000101010000101010010101
11001100010100011011000010
01100010100110010110001101

DESIGN CONFERENCE

TELL ME MORE about the Analog & Mixed-Signal Design Conference.

I'm interested in ☐ Attending ☐ Exhibiting

Name

Title


Company

Street


City, State, Zip

Phone () Fax ()

68040 + '030 I/O Bus = 39 MIPS



Available: 4-128MB RAM, Ethernet, SCSI, VSB, VME64



Taurus™ Dual Bus Architecture

The Taurus is a dual-processor, dual-bus, single slot VME board. Its dual-bus architecture allows the 68040 to execute code uninterrupted, while the '030 processes on-board I/O. This optimizes the 68040's performance. Using the '030 as an I/O processor simplifies writing your code. You only need to write high level code to the 68040. The '030 handles the device level code. You also can use the '030 as a DMA controller, while the 68040 directly controls all on-board I/O devices. The '030 uses the SRAM with the 128KB of EPROM code provided by Omnibyte.

| | |
|-----------------|---|
| Performance | 68040: 29 MIPS, '030: 10 MIPS, VME: 50MB/sec, VSB†: 50MB/sec |
| Intelligent I/O | Ethernet: i82596CA†, SCSI: NCR53C710†, 4 RS232D: CD2401 |
| Standard I/O | 2 RS232D: 68C681 DUART, 32 Lines Parallel I/O, or 16 w/ Centronics Printer Port |
| Memory | 4MB to 128MB† DRAM, 512KB SRAM†, 8KB NVRAM, 1MB FEPROM†, 4MB EPROM |
| Other | VSB†, VME64†, Watchdog, Calendar Clock, Mailbox, (6) 16-bit Timers, Snooping, Advanced Omnimodule™ Socket |
| Software | VxWorks ¹ , OS-9 ² , UNIX ³ CrossCodeC, FreeForm ⁴ , OMNIBug |

† Denotes optional features.

The Taurus extensively uses intelligent, on-chip DMA devices for Ethernet, SCSI and serial I/O. This helps reduce processor intervention. Up to 2 stackable modules contain the DRAM. This allows upgradable options from 4-128MB.

Advanced Omnimodules provide additional custom I/O. You can stack Advanced Omnimodules up to 3 high. The Taurus can accept 1 memory module and 1 Advanced Omnimodule and still fit into a single slot.

To learn more contact Larry Snow:

800-638-5022


Circle No.120

© 1992 Omnibyte Corp.

OMNIBYTE™ Corporation, 245 W. Roosevelt Rd., West Chicago, IL 60185

Trademarks: 1-Wind River Systems, 2-Microware Systems Corp., 3-UNIX Systems Laboratories, Inc., 4-Software Development Systems Inc.

Searching for a R3000 VME Hardware/Software Solution?



Omnibyte has the most complete line of MIPS R3000 VME hardware, software and systems for your project. In fact you may select from 72 distinct VME configurations.

You can get the fast and powerful Pulsar 3000™, shown above, featuring:

- 25MHz R3000A RISC CPU
- 25MHz R3010A Floating-Point Coprocessor (opt.)
- 4 R3020 Write Buffers
- 128KB (or 32KB) I-cache
- 128KB (or 32KB) D-cache
- 32, 16, 8 or 4MB DRAM
- High speed SCSI port (opt.)
- Ethernet interface (opt.)
- 4 Serial ports
- 4 ROM sockets (up to 4MB)

- Real time calendar clock w/battery
- 2KB NV RAM
- VIC068 VME Interface (slot 1) Controller

And the Pulsar is available with the following software:

| Title | Type |
|--|-------------------------------------|
| SPP/e ¹ | PROM monitor debugger & I/O library |
| VADS ² | ADA compiler |
| VxWorks ³ | Real time o.s. |
| UniPlus ⁴ SVR4 ⁵ | UNIX ⁵ o.s. |

In addition to our boards and software, we can also provide you with a complete line of development systems.

For further information, give Larry Snow a call today at

800-638-5022


CIRCLE NO. 119

OMNIBYTETM CORPORATION
245 W. Roosevelt Rd.
West Chicago IL 60185

Trademarks: 1-MIPS Computer Systems Inc. (available in binary or source), 2-VERDIX Corporation, 3-Wind River Systems (available for MIPS & DEC workstations), 4-UniSoft Group Limited, 5-UNIX Systems Labs

High-speed A-D converters shift to new architectures

Jeffrey Child, Senior Editor



Hybrid A-D converters such as Comlinear's CLC935 still have a performance advantage over monolithic parts. This 12-bit, 15-Msample/s part maintains a 77-dB spurious-free dynamic range at 7.22 MHz. Signal-to-noise ratio is 66 dB out to Nyquist rates (7.5 MHz).

Driven by the demands of applications such as professional video, radar, high-definition television, and ultrasound, designers have an endless appetite for faster analog-to-digital conversion. And with digital signal processing becoming commonplace in today's systems, they need converters that offer more than just fast digitizing of an analog signal. More than ever, data coming out of a converter is being manipulated, stored, compressed, or otherwise changed using DSP techniques. As a result, designers are demanding parts that also boast high resolution and better dynamic specifications.

Responding to these demands, makers of high-speed A-D converters are striving to offer high resolution at high speed. At the same time, some vendors have created monolithic parts with performance nearly equal to that of more expensive hybrid converters. To accomplish these feats, makers of hybrid and monolithic parts are squeezing all they can out of traditional converter architectures and process technologies.

A few have seen the writing on the wall, however, and are abandoning classical A-D converter architectures such as two-pass flash and successive-approximation in favor of new proprietary schemes. "There's a whole movement going on in architecture, design and technology, but buried underneath is the issue of performance, because systems often live and die on how well the con-

verters work," says Alan Hansford, product marketing manager for the data conversion group at Comlinear (Fort Collins, CO).

■ The same for less

As A-D converter vendors push the performance envelope, they're also pushing down cost. A little over a year ago, a converter with a speed of 10 Msamples/s at 10-bit resolution was only available as a hybrid or a board-level product that cost up to \$1,000. Today, monolithic parts are available with the same performance but at a fraction of the cost. For some applications this trend has opened up new choices for how systems are designed.

Exemplifying the drive toward high performance at a lower cost is the SPT792X, a family of A-D converters from Signal Processing Technologies (SPT—Colorado Springs, CO). Available with both 10- and 30-Msample/s speeds, these parts are the industry's first 12-bit, high-speed, monolithic A-D converters. Coming in ECL- and TTL-output configurations, SPT792X converters typically consume less than 1.1 W of power. The 30-Msample/s version offers a signal-to-noise ratio (SNR) of 67 dB (typical).

Offering this performance and speed on a single die forced SPT to break away from traditional architectures. Although the company wouldn't reveal any details of its patent-pending architecture, Rick Mintle, director of marketing, hints

that the implementation of the on-chip sample-and-hold amplifier (SHA) was a key element. According to Mintle, the architecture let SPT use 80 percent less power than a hybrid at 1/5 the cost. SPT also plans to sell the parts in die form, letting you install them on multichip modules.

■ An A-D converter per channel

The availability of low-cost, high-performance data converters lets you make architectural changes to your system—adding more channels, for example. "A year or two ago, anybody using a 12-bit, 10-MHz A-D converter would only use one or two per system because they cost anywhere from \$700 to \$1,000 each," says Mintle. "And even the best hybrids were very power hungry—around 5 to 7 W. Now you can get fast converters for \$150 each. For the same price, then, you can have eight times as many converters. Especially in the higher volume applications, this opens up new system architectures that were once prohibitive because of cost. Now, for example, you can have an A-D converter per channel, where before you had to multiplex several channels through one expensive, high-performance ADC." Also, since any extra component in the circuit path causes delays, eliminating a multiplexer improves overall system performance.

For its part, market leader Analog Devices (Wilmington, MA) also has a 12-bit, 10-Msample/s part, the AD872, in the works. The AD872,

| Model | Sample rate (Msample) | Resolution (bits) | Input bandwidth (MHz) | Linearity ϵ (\pm LSB) | Power dissipation (W) | Sample & hold | Internal reference | Package | Package technology | Price | Comments |
|---|-----------------------|-------------------|-----------------------|-----------------------------------|-----------------------|---------------|--------------------|--------------------------|--------------------|--------------------------|--|
| Analog Devices Three Technology Way, PO Box 9106, Norwood, MA 02062 (617) 329-4700 Circle 301 | | | | | | | | | | | |
| AD773 | 18 | 10 | 100 | 0.75 | 1.2 | yes | no | 28-pin DIP | monolithic | \$55 (100s) | pipelined architecture; TTL compatible |
| AD872 | 10 | 12 | 80 | 2.5 | 1.15 | yes | yes | 28-pin DIP; 44-pin LCC | monolithic | \$165 (100s) | pipelined architecture; TTL compatible |
| AD9005A | 10 | 12 | 38 | 1 | 3.2 | yes | yes | 46-pin DIP | hybrid | \$575 (100s) | subranging flash; TTL compatible |
| AD9014 | 10 | 14 | 60 | 0.75 | 12.8 | yes | yes | PCB | hybrid | \$2750 (100s) | flash; ECL compatible |
| AD9020 | 60 | 10 | 175 | 1.25 | 2.8 | no | no | 68-pin LCC | monolithic | \$165 (100s) | flash; TTL compatible |
| AD9032 | 25 | 12 | 150 | 0.7 | 5.4 | yes | yes | 40-pin DIP | hybrid | \$185 (100s) | subranging flash/hybrid; ECL compatible |
| AD9034 | 20 | 12 | 150 | 1 | 4.5 | yes | yes | 40-pin DIP | hybrid | \$715 (100s) | subranging flash; ECL compatible |
| AD9060 | 75 | 10 | 175 | 1 | 2.8 | no | no | 68-pin LCC | monolithic | \$185 (100s) | flash; ECL compatible |
| Burr-Brown 6730 S Tucson Blvd, Tucson, AZ 85706 (602) 746-1111 Circle 302 | | | | | | | | | | | |
| ADS605 | 10 | 12 | 15 | 0.9 | 1.75 | yes | yes | 40-pin 0.6-in. wide DIP | hybrid | \$250 (100s) | — |
| ADC603 | 10 | 12 | 30 | 0.9 | 6.1 | yes | yes | 46-pin 1.56-in. wide DIP | hybrid | \$425 (100s) | — |
| Comlinear 4800 Wheaton Dr, Fort Collins, CO 80525 (303) 226-0500 Circle 303 | | | | | | | | | | | |
| CLC925 | 10 | 12 | 70 | 0.35 | 4.2 | yes | yes | 40-pin side-brazed DIP | hybrid | \$449 (100s) | 74-dBc spurious-free dynamic range; 67-dB S/N ratio; subranging converter |
| CLC935 | 15 | 12 | 135 | 0.35 | 5.2 | yes | yes | 40-pin side-brazed DIP | hybrid | \$650 (100s) | 77-dBc spurious-free dynamic range; 67-dB S/N ratio; subranging converter |
| CLC936 | 20 | 12 | 135 | 0.35 | 5.2 | yes | yes | 40-pin side-brazed DIP | hybrid | \$750 (100s) | 73-dBc spurious-free dynamic range; 65-dB S/N ratio; complete subranging converter |
| Datel 11 Cabot Blvd, Mansfield, MA 02048 (508) 339-3000 Circle 304 | | | | | | | | | | | |
| ADS-119 | 10 | 12 | 50 | 1 | 1 | yes | yes | 24-pin DDIP | hybrid | \$487 (OEM) | — |
| ADS-120 | 20 | 12 | 60 | 1 | 3.8 | yes | yes | 40-pin TDIP | hybrid | \$790 (OEM) | TTL-compatible output |
| ADS-945 | 10 | 14 | 20 | 0.75 | 3.5 | yes | yes | 40-pin TDIP | hybrid | \$790 (OEM) | — |
| ILC Data Device 105 Wilbur Pl, Bohemia, NY 11725 (516) 567-5600 Circle 305 | | | | | | | | | | | |
| ADC-00110 | 10 | 12 | 50 | 2 | 8 | yes | yes | 46-pin DIP | hybrid | \$1200 (1-4 pcs) | — |
| Micro Power Systems 3100 Alfred St, Santa Clara, CA 95054 (408) 562-3670 Circle 306 | | | | | | | | | | | |
| MP7695 | 1 | 10 | 20 kHz | 1 | 50 mW | yes | no | 24-pin PDIP, CDIP, SOIC | monolithic | \$10.79 - \$86.50 (100s) | — |
| MP8782 | 5 | 10 | 5 | 0.5 | 200 mW | yes | yes | 24-pin PDIP, CDIP, SOIC | monolithic | — | — |
| MP8790 | 2 | 12 | 2 | 2 | 250 mW | yes | no | 24 & 28 SOIC; 52 PLCC | monolithic | — | — |

| Model | Sample rate (Msample/s) | Resolution (bits) | Input band- width (MHz) | Linearity error (\pm LSB) | Power dissipa- tion (W) | Sample & hold | Internal reference | Package | Package technology | Price | Comments |
|---|----------------------------|----------------------|----------------------------|---------------------------------|----------------------------|---------------|-----------------------|---|-----------------------|--------------------------|--|
| Signal Processing Technologies 1510 Quail Lake Loop, Colorado Springs, CO 80906 (719) 540-3900 Circle 307 | | | | | | | | | | | |
| SPT7810 | 20 | 10 | 120 | 1 | 1.3 | yes | no | 28-pin DIP | monolithic | \$54 - \$64 (100s) | ECL logic; 59-dB S/N ratio |
| SPT7814 | 40 | 10 | 120 | 1 | 1.3 | yes | no | 28-pin DIP | monolithic | \$81.10 - \$95.90 (100s) | ECL logic; 57-dB S/N ratio |
| SPT7820 | 20 | 10 | 120 | 1 | 1 | yes | no | 28-pin DIP | monolithic | \$54 - \$64 (100s) | TTL logic; 59-dB S/N ratio |
| SPT7824 | 40 | 10 | 120 | 1 | 1 | yes | no | 28-pin DIP | monolithic | \$81.10 - \$81.10 (100s) | TTL logic; 57-dB S/N ratio |
| SPT7910 | 10 | 12 | 120 | 2 | 1.4 | yes | no | 32-pin DIP | monolithic | \$150 (100s) | ECL logic; 67-dB S/N ratio |
| SPT7920 | 10 | 12 | 120 | 2 | 1.1 | yes | no | 32-pin DIP | monolithic | \$150 (100s) | TTL logic; 67-dB S/N ratio |
| SPT7922 | 30 | 12 | 120 | 2 | 1.1 | yes | no | 32-pin DIP | monolithic | \$250 (100s) | TTL logic; 66-dB S/N ratio |
| Sipex 22 Linnell Cir, Billerica, MA 01821 (508) 667-8700 Circle 308 | | | | | | | | | | | |
| SP9560 | 10 | 12 | 10 | 0.5 | 3.25 | yes | yes | 46-pin DIP | hybrid | \$700 (100s) | — |
| Sony Component Products 10833 Valley View St, Cypress, CA 90630 (800) 288-SONY Circle 309 | | | | | | | | | | | |
| CXA1496AQ | 20 | 10 | 20 | 1 | 0.31 | no | no | 48-pin QFP | monolithic | \$93 (100s) | subranging A-D converter |
| TRW LSI Products 4243 Campus Point Ct, San Diego, CA 92121 (619) 457-1000 Circle 310 | | | | | | | | | | | |
| TDC1020 | 20 | 10 | 5 | 0.5 | 7.2 | no | no | 64-pin CDIP; 68-pin PGA | monolithic | \$93 (100s) | flash A-D converter |
| TDC1049 | 30 | 9 | 15 | 0.5 | 5.22 | no | no | 64-pin CDIP; 68-pin PGA; 68-pin LCC | monolithic | \$90 (100s) | DESC SMD 5962-88532 available for military flash A-D converter |

IEE BOOKS

Circuits & Systems Series

Analogue-digital ASICs: circuit techniques, design tools and applications **\$94.00**

R.S. Soin, F. Malaberti and J. Franca (Eds.)
This book comprises eighteen chapters arranged in four sections: Processing Technology; Circuit Techniques and Building Blocks; Design Applications; and CAD Supporting Tools.
431pp., 229X148mm, Casebound
ISBN 0 86341 259 9, 1991 Circuits and Systems Series #3



Analogue IC Design: the current-mode approach **\$94.00**

C. Toumazou, F.J. Lidgey and D.H. Haigh (Eds.) (Joint winner of the 1991 Rayleigh Book Award)
With this book the editors have produced a "Bible" for modern integrated circuit design techniques. It is an excellent complement to the first volume in the Circuits and Systems series and will be of benefit to students, both undergraduate and postgraduate, and researchers and circuit designers in academia and the Electronics Industry
646pp., 229X148mm, Casebound
ISBN 0 86341 215 7, 1990 Circuits and Systems Series #2

GaAs Technology and its impact on circuits and systems **\$90.00**

D.G. Haigh and J. Everard (Eds.)
This book aims to lay down sound principles for forthcoming developments in which both the speed and the level of integration complexity of GaAs circuits will increase dramatically, meeting needs in such areas as ultra-fast computers, personal communications and new broadcasting technology.
457pp., 229X148mm, Casebound
ISBN 0 86341 187 8, 1989 Circuits and Systems Series #1

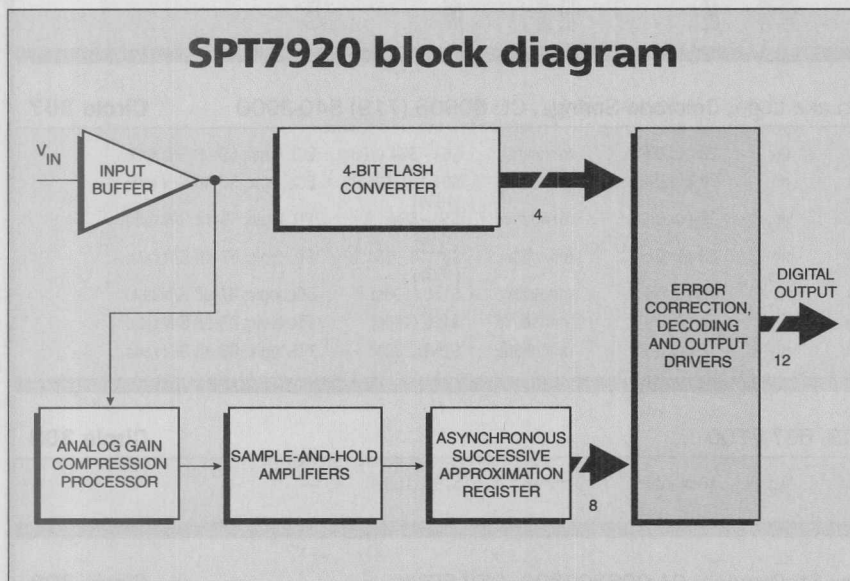
INSPEC/PPL, IEEE Service Center

445 Hoes Lane, Piscataway, NJ 08855-1331 Phone: (908) 562-5553 • FAX: (908) 981-0027

CIRCLE NO. 72

INTEGRATED CIRCUITS

SPT7920 block diagram



The SPT7920 from Signal Processing Technologies is the industry's first 12-bit, high-speed, monolithic A-D converter. Available with 10-Msample/s and 30-Msample/s speeds, the chip offers a signal-to-noise ratio of 67 dB (typical). To achieve this level of performance, SPT used a non-traditional, proprietary architecture.

like SPT's device, uses a non-traditional architecture. It's based on flash but uses a multistage pipeline that reduces the number of comparators needed. This architecture was first used in the AD773, Analog Devices' 10-bit, 18-Msample/s monolithic converter.


"Using that multistep pipelined architecture results in a dramatic reduction of on-chip comparators—from 1,000 to 64," says Paul Errico, video and imaging marketing manager at Analog Devices. In using fewer comparators, the AD773 dissipates 1.2 W (typical), much less than conventional flash converters. Analog Devices used the AD773 core as a basis for other parts, including the AD872.

Addressing the demands for both lower cost and higher speed, Analog's AD773 A-D converter, at \$55, has the distinction of being the lowest-cost 10-bit, 18-Msample/s device on the market. The AD773 further reduces system cost by eliminating the need for external buffers and ECL-to-TTL interface circuitry. A front-end SHA

REAL-TIME

SPARC

Integrated VME Workstation



*The New AP Labs VMEstation™:
Fully integrated systems for seamless real-time development...*

- ▶ SPARCengine or 68000 series architecture
- ▶ SunOS and VxWorks with AP Labs' Asynchronous I/O
- ▶ RDBXworks™: C source-level SPARC debugger
- ▶ Fully integrated VME and SBus I/O real-time device driver support
- ▶ Data Acquisition; Communications; Signal Processing

AP Labs

6215 Ferris Square
San Diego, CA 92121
(619) 546-8626
Fax: (619) 546-0278

SPARC, SPARCengine, SunOS and SBus are trademarks of Sun Microsystems, Inc. VxWorks is a trademark of Wind River Systems. 68000 is a trademark of Motorola, Inc.

CIRCLE NO. 73

"FUZZY LOGIC" REPRINTS

Last year our March special report on **Fuzzy Logic** generated over **4,100 reader inquiries** and the **April 1992 Fuzzy Logic report** is already poised to generate an even higher number of inquiries! **Reprints of the April article are available now**—quantities of 100 or less are \$1.00 each, quantities over 100 are .75 each. To order, call Patti Kenney at (508) 392-2124.

COMPUTER DESIGN Technology and Design Directions

CIRCLE NO. 74

permits signal digitization beyond the sample rate. The part has a typical SNR-plus-distortion of 56 dB for 1-MHz sine waves.

■ Process innovations next

Architectures aren't the only things to change in the A-D converter world. Most vendors claim to have new process technologies in the works that will be key for their next jumps in performance. "We have a new process, but at the moment we're not doing any high-speed converters on it," says Analog Devices' Errico. "When we do, the flash converters and even the pipelined architectures will probably drop an order of magnitude in power consumption. Once again, you run into a trade-off—if you want it to be faster, it's going to eat up more power. It'll take a couple of years for most companies to go to their new processes. They'll start by putting their big-money chips, like op-amps, on the new processes first."

Comlinear, a maker of hybrid A-D converters, recognizes that innovation in process technology is critical. "Now we're seeing monolithics coming in with lower cost and lower power, but also lower performance, than the hybrids. Meanwhile, the traditional subranging hybrids are stalling in terms of technology and performance," says Comlinear's Hansford. "We have a 12-bit, 30-Msample/s hybrid in development now that's due out in January. As far as we know, that part is *it* for subranging hybrids. We don't know how to make it go any faster without a real dramatic change in process technology."

According to Hansford, Comlinear has invested a lot of money in a new advanced architecture, one that's

radically different from both Analog Devices' and SPT's. "With the new architecture, we believe that we can start anew from that 12-bit, 30-Msample/s performance roadblock and go significantly faster and to higher resolution right away," says Hansford. "We have working chips in-house now that indicate that the architecture functions."

Comlinear's hottest A-D converters that are available now are the CLC935 and CLC936, a pair of 12-bit hybrids with speeds of 15 Msamples/s and 20 Msamples/s respectively. Optimized for dynamic specifications, the 12-bit, 20-Msample/s CLC935 features a typical value of 77-dB spurious-free dynamic range. This value is maintained over a wide frequency range. The part's designed for communications receivers, high-dynamic-range analyzers (such as spectrum analyzers) and radar and guidance. It has an elaborate SHA that determines the distortion products of the converter. Although it's targeted for frequency-domain applications, its SNR of 66 dB makes it appropriate for time-domain applications as well.

In the 20-Msample/s CLC936, dynamic range and performance were traded off to get that extra 5 Msamples/s of speed. Yet despite the faster speed of the CLC936, Hansford candidly admits that the CLC935 is a better part, since he stresses the importance of dynamic specifications in today's system design. "It's an expansion of the performance envelope, whereas the CLC936 is just more of the same. These are subtle topics that can take a system designer a long time to find out about." ■


We want to hear from you!

So, what do you think of Computer Design? What do you like best, the Technology Viewpoint pieces? The Technology Directions articles? The Special Reports? The Product Focus articles and specification tables? New Product Developments? Steve Ohr's column on Mixed-Signal Design? What is there about *Computer Design* that you don't like? What would you like to see us cover more extensively?

The only way we can know what you like or don't like, or what you want to see more of or less of, is if you'll write us. There's about 11.25 in.² of space on the Reader Inquiry Card for your comments, but if that's not enough, you can write directly to John Miklosz, Editor-in-Chief, *Computer Design* magazine, 1 Technology Park Dr, Westford, MA 01886.

Think serial.

Communicate in your choice of protocols using IndustryPack™ I/O modules. IP-based serial communication gives you more choices in fewer slots for less money.



Choose from
25 IndustryPacks
and seven carrier

boards to create the single-slot solution that meets your communication requirements.

Pack up to 32 lines in one slot. Whether you need asynch, synch, or T1. Interface RS-232, RS-422, RS-485, non-standard or even optically isolated — IPs support them all.

And single-slot solutions start at under \$48 per line for VMEbus, Nubus and ATbus.

Best of all, 68030-based intelligent carriers give you the highest performance, most versatile communication boards available anywhere.

If you're ready to think serial, put your call to GreenSpring at the head of the queue to receive a free catalog. Or fax us your requirements today.

- IP-Serial
- IP-Octal 232
- IP-Octal 422
- IP-Octal Opto Serial

GREEN SPRING
COMPUTERS

1204 O'Brien Drive, Menlo Park, CA 94025
(415)327-1200/FAX(415)327-3808

THE CHOICE WITH A FUTURE

CIRCLE NO. 75

Register-rich FPGA family has 9-ns interconnect delays

Implemented in SRAM technology, the FLEX 8000 FPGA family from Altera features in-circuit reconfigurability and low standby power, as well as a continuous global interconnect structure that results in 9-ns across-the-chip interconnect delays. Unlike Altera's Classic and MAX parts, which are optimized for combinatorial logic, FLEX—which stands for Flexible Logic Element Matrix—devices are register-intensive, with flip-flop counts of from 452 to 2,252 and usable gate counts ranging from 4,000 to 24,000. The FLEX architecture provides two levels of logic granularity: small, fine-grained logic elements and larger, coarse-grained logic array blocks. The in-system clock rate is 70 MHz.

LUTs and flip-flops

With a 5-ns propagation delay, each logic element contains a four-input look-up table (LUT) and a configurable flip-flop. The LUT can synthesize any combinatorial function of four variables. You can configure the flip-flop for D, T, JK, or SR operation.

A logic array block (LAB) is made up of eight logic elements fully interconnected through a local interconnect. The TRUE and COMPLEMENT of four fast, low-skew lines driven from device input pins enter each LAB for global control signals. Two may be used for register clocks and two to clear registers, or one of the CLEAR lines may be used to preset a register. The clock and clear control signals for each logic element's flip-flop can come from fast input pins, any other device pin or any internally generated signal.

Called FastTrack, this row-and-column interconnect scheme in the FLEX devices consists of dedicated metal lines that transport signals in both horizontal and vertical direc-

tions with three types of fixed interconnect delays. Any two logic elements can be connected through these continuous metal lines.

Two high-speed datapath chains bypass this logic element intercon-

nect structure.

At the end of each FastTrack routing channel are I/O elements that provide bidirectional I/O buffers with output slew-rate control and a flip-flop for either an input or output register.

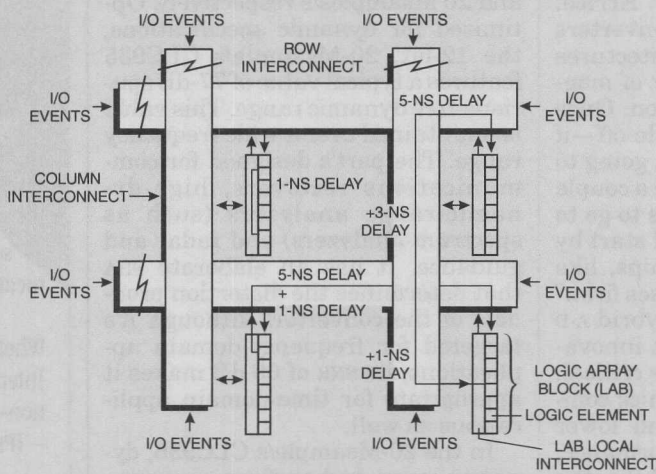
You can download your FLEX device configurations from EPROM, EEPROM or system RAM, depending on the target application. FLEX devices can be dynamically reconfigured on a personal computer board in less than 100 ms.

Members of the FLEX 8000 family will be available in versions that support the IEEE-1149.1 JTAG boundary test scan standard. Altera's MAX+Plus II design tools, available in both Windows and UNIX versions, support FLEX devices.

The EPF81188 device is the first available FLEX family member. It has 12,000 usable gates (24,000 total), 1,188 flip-flops and up to 232 pins. It will be available in a 232-pin PGA for \$495 each in 100-piece quantities. Sampling will begin this quarter, with production quantities in mid-1993.

—Barbara Tuck

FastTrack interconnect scheme



The FastTrack continuous row-and-column interconnect scheme in Altera's new FLEX FPGA family provides fixed interconnect delays. All logic elements within the same logic array block (LAB) are interconnected by a local interconnect with a 1-ns delay. Any two logic elements in the same row have a 6-ns interconnect delay (LAB local interconnect + row interconnect), and any two logic elements in different rows have a 9-ns interconnect delay (LAB local interconnect + row interconnect + column interconnect).

nect structure. Called cascade and carry chains, they run from logic element to logic element along each row of the device to support the high-speed operation of counters, adders and logic functions with wide inputs. To implement logic functions with wide inputs without the cumulative delays of serially connecting entire logic elements, a cascade chain can connect LUT outputs of adjacent logic elements in parallel.

Carry and cascade

The LUT also generates sum and carry functions for a carry chain. The carry chain, which connects adjacent logic elements in the same way that the cascade chain does, lets you build fast counters and adders. Altera claims that FLEX interconnect delays are typically an order of mag-

FLEX 8000 at a glance

- Based on SRAM technology
- Five devices ranging from 4,000 to 24,000 usable gates
- 9-ns across-the-chip interconnect delay
- Optimized for register logic
- Both fine and coarse levels of logic granularity
- I/O registers on all I/O pins
- 70-MHz in-system clock rate

Altera Semiconductor

2610 Orchard Pkwy
San Jose, CA 95134-2020
(408) 984-2800

Circle 355

DSP chip set aimed at real time

A digital signal processor chip set consisting of a DSP and an address generator has been designed by Sharp Microelectronics to provide real-time response for DSP applications. Billed as the world's fastest DSP chip set, the LH9124 processor and LH9320 address generator pack a lot of performance into two compact packages. The heart of the chip set, the DSP itself, is a 24-bit, fixed-point processor available in 33- and 40-MHz versions.

Fabricated using Sharp's 0.8- μ m, double-metal technology, the processor boasts computational capabilities such as four 24-bit complex datapaths, six 24-bit multipliers and dual 60-bit accumulators. In addition, the chip includes 26 high-level opcodes, among them real and complex Radix-2, Radix-4, Radix-16 butterfly operations, and real and complex FIR filter operations. It directly performs discrete and fast Fourier transforms, discrete and fast cosine transforms, convolutions, correlations, multirate arbitrations, and multichannel buffering.

Multiport architecture

The LH9124 has two bidirectional data ports, a bidirectional acquisition port and a bidirectional coefficient port. The 24-bit-wide multiport data-flow structure makes external multiplexing of data unnecessary and enables higher throughput and reduced component count. The architecture lets each bus be served by simple off-the-shelf, single-port memories.

To handle a wide range of values, the processor uses block floating-point. Scaling can be performed automatically, or you can supply your own scaling factors. Internal bus structures carry full precision all through the chip.

The bidirectionality of the design also lends itself to the development of recursive systems. The six on-board multiplier/accumulators let the processor outperform single-multiply/accumulate DSPs.

The address generator includes over 150 embedded memory address sequences, synchronized to dramatically reduce the software investment required. With its 20-bit address, the LH9320 address

generator can directly address a full megaword of memory, allowing for very large arrays, 2-D arrays or support for up to 32 independent channels.

The combination of on-chip resources and built-in opcodes results in some impressive performance. For example, Radix-16 butterflies can be implemented in 400 ns and a 1K complex FFT in only 80 ns. In addition, the architecture lends itself to cascading multiple 9124s; three can complete a complex FFT in 25 ns.

Typical systems comprise a single DSP chip and three address generators, each supporting corresponding memories. The evaluation board offered by Sharp supplies up to 500 MFlops of processing on a small platform. Software development support is provided by the LH9124 and LH9320 Real-Time Simulation Kits. Hardware development is supported by the chip-set evaluation module and user guide.

The price for the 40-MHz processor is \$1,200, and the address generator is \$250. Both are sold in 100-piece lots. The real-time simulators sell for \$300 each, and the evaluation module is available for \$5,000.

— Warren Andrews

DSP chip set at a glance

- 33- or 40-MHz DSP chip and address generator
- 262-lead PGA processor, 68-lead PLCC address generator
- 26 built-in high-level functions
- Dual 60-bit accumulators
- Over 150 memory address sequences
- On-chip block floating-point with automatic or user scaling factors
- Expanding internal bus structure carries full precision

Sharp Microelectronics
5700 NW Pacific Rim Blvd #20
Camas, WA 98607
(206) 834-2500

Circle 351


Think parallel.

Digital I/O IndustryPacks™ give you more ways to connect more lines in less space for less

money. Up to 192 parallel lines fit in one VME or ATbus slot – up to 96 lines in a Nubus slot.

Whether you need optical isolation, high voltage/high current drive, interrupts on every input or simply lots of reliable, inexpensive I/Os, chances are that one of our six convenient digital IndustryPacks will meet your requirements.

All IPs work in all slots, so you can easily create the interface solution you need from over thirty IndustryPacks and IP carrier boards.




If saving space and money has you thinking parallel, call or fax GreenSpring Computers for a free catalog today.

- IP-Digital 24
- IP-Digital 48
- IP-Dual PI/T
- IP-ADIO
- IP-Opto Interrupter
- IP-Opto Driver

GREEN SPRING
COMPUTERS

1204 O'Brien Drive, Menlo Park, CA 94025
(415) 327-1200/FAX (415) 327-3808

THE CHOICE WITH A FUTURE



YOU ARE CORDIALLY INVITED TO THE...

Computer Peripherals Invitational Computer Conference

featuring new technology presentations and product displays

Request your *Complimentary Invitation* and Agenda

OEMs, Systems Integrators, VARs — If you are involved in the purchasing, design, or specification of peripheral products... you should be here!

Copy your business card here
FAX to: (714) 476-9969


a division of
Dataquest®
 a company of
The Dun & Bradstreet Corporation
Dataquest is a registered trademark of A.C. Nielsen Company

Indicate your conference selection below
and return this section by FAX or mail.

- | | |
|--|------------|
| <input type="checkbox"/> Newton, MA | Sep 09 '92 |
| <input type="checkbox"/> Dallas, TX | Sep 22 '92 |
| <input type="checkbox"/> Portland, OR | Oct 22 '92 |
| <input type="checkbox"/> Irvine, CA | Jan 07 '93 |
| <input type="checkbox"/> San Jose, CA | Feb 09 '93 |
| <input type="checkbox"/> Minneapolis, MN | Mar 04 '93 |
| <input type="checkbox"/> Austin, TX | Mar 25 '93 |
| <input type="checkbox"/> Denver, CO | Apr 08 '93 |
| <input type="checkbox"/> Nashua, NH | Apr 20 '93 |

3990 Westerly Place, Suite 100, Newport Beach, CA 92660 Tel: (714) 476-9117 Fax: (714) 476-9969

IPs get PCMCIA adapter

Memory upgrade solutions to industrial automation problems are often far more complex and costly than their counterparts in the PC environment. X. Kim Rubin, vice-president of GreenSpring Computers, says, however, that "PCMCIA-compatible memory prices are often far below distributor prices for discrete chips and, in at least one case, flash memory has been quoted lower in PCMCIA packages than as discrete chips, and in smaller volumes."

To capitalize on this market for laptop-compatible memory cards, GreenSpring has added a PCMCIA-to-IndustryPack (IP) adapter to its family of standard IP I/O modules. Dubbed the IP-PC Card, the module provides access to PCMCIA memory cards by offering a complete PCMCIA interface to GreenSpring's IP interface.

The module is a two-part system with an IndustryPack adapter card attached to a PCMCIA receiver card by a flexible circuit. "As of this time," says Rubin, "the maximum distance between the IndustryPack site and the adapter card has been targeted at 24 in., but most likely, initial versions will have a 16-in. connector."

The rationale behind separate locations for the two parts is that the PCMCIA adapter can be located away from the IP carrier to let PCMCIA cards be user-accessible in embedded applications. This can be important in applications such as medical instrumentation and various data logging and equipment monitoring chores. The PCMCIA cards are far more rugged than rotating media, and already a broad variety of connectors, including dust- and contamination-resistant varieties, are available.

While Rubin is confident that compatibility problems will be resolved, the PCMCIA community's manufacturers are implementing slightly different versions of the standard. And while the changes made to Version 2.0 of the PCMCIA standard corrected most of its hardware problems, software definitions are still wanting.

On the positive side of the ledger, however, many of the memory interface compatibility issues facing PCMCIA are exclusive to DOS. "In our implementation," says Rubin, "we

simply map the entire memory space into the host machine in a free area of the memory. We don't have to worry about mapping the memory to the architecture of the machine with register windows, as happens in a DOS environment."

The memory has to look like it's in a DOS file structure for it to work in Intel-based machines, and this complicates the situation some. In most current applications, IP-based systems operate in real time using one of the popular real-time operating systems (RTOS). "We already provide a flash-memory file system under OS-9, and it's not difficult to operate the PCMCIA with other RTOS," Rubin adds.

Despite the confusion about the standard, there's a lot of momentum behind the PCMCIA specification. Not only is it being pushed from the chip side, but several large Japanese consumer companies, among them Panasonic and Sony, have elected to back it as a laptop standard.

Already flash memory from one manufacturer that's available at \$46/Mbyte in discrete chips in 10,000-piece order quantities can be bought from the same vendor at \$36/Mbyte in PCMCIA packages for 1,000-piece orders.

GreenSpring's IP-to-PCMCIA adapter will be available in the fourth quarter at \$295. Initial versions will be equipped with a flex circuit between the two modules. Other versions are being planned that will use a conventional ribbon cable.

— Warren Andrews

IP-PC Card at a glance

- IndustryPack-to-PCMCIA adapter
- Connects to any IndustryPack site
- Uses inexpensive PCMCIA connector
- 16-in. cable between PCMCIA and IndustryPack
- Uses inexpensive flash memory


GreenSpring Computers

1204 O'Brien Dr
Menlo Park, CA 94025
(415) 327-1200

Circle 352

Think nonlinear?

Nonlinear thinking helps solve problems, but... if your requirements include data acquisition or signal processing, you need one of our highly linear A/D or D/A converters.



Fit up to 80 analog channels into one VMEbus or ATbus slot — up to 40 per Nubus slot. Directly connect most sensors including thermocouples and current loops. Programmable gain and low drift guarantee precision data.

If you need 16-bit DACs or 16 separate DACs, choose from five precision ADC and DAC IndustryPacks™ to build your analog subsystem in less space with less heat, drift, software or expense than alternative solutions.

If you're ready for linear solutions to your data acquisition or signal processing needs, call or fax us your requirements today. Think efficiency... think GreenSpring!

- IP-ADC
- IP-DAC
- IP-ADIO
- IP-DenseDAC
- IP-16DAC

GREEN SPRING
COMPUTERS

1204 O'Brien Drive, Menlo Park, CA 94025
(415) 327-1200/FAX (415) 327-3808

THE CHOICE WITH A FUTURE

EDA suite boasts front-to-back tool integration

With the introduction of its System Workbench, Cadence Design Systems has combined its EDA products with the Allegro Correct by Design (CBD) tools it acquired in the merger with Valid Logic Systems at the end of 1991. System Workbench integrates CBD

ical implementation of your design doesn't violate any specified electrical or physical rules.

System Workbench supports technology independence across all of its tools to provide design retargetability and optimization. During design entry using the Composer tool, for

mission effects such as crosstalk, can be back-annotated automatically from the physical layout to logic or timing simulations to verify circuit performance.

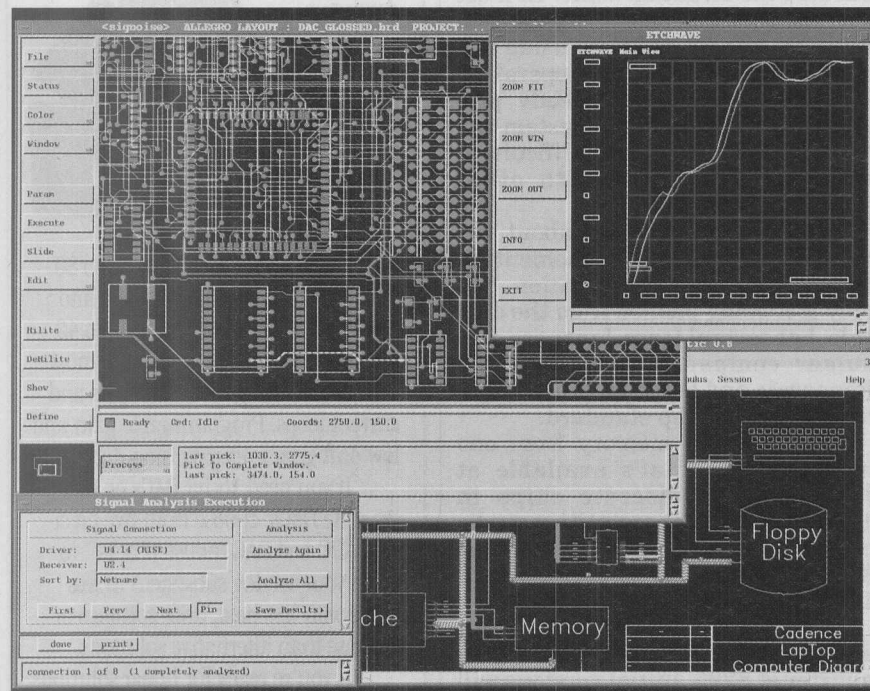
Warp-4 maps space

A new autorouter, Prance-XL, has also been bundled with the system. A key element of the router is a technology that Cadence calls Warp-4; it maps available space into contours that flow around pads and other obstacles to achieve maximum trace density. These traces appear curved on the screen but the router sees them as straight, letting it proceed at full speed. The Warp-4 technology automatically adheres to manufacturing considerations, such as centering traces, as it flows between pads.

The Prance-XL/Allegro CBD integration uses a router backplane to combine CBD and Warp-4 technologies. The autorouter automatically adheres to all physical design rules and constraints that have been established within the System Workbench environment. This lets you define critical nets as "no rip-up" or "no-shove," with Prance-XL adhering to these constraints.

System Workbench is available now, priced from \$58,000 to \$145,000, depending on the configuration.

—Mike Donlin



Cadence's System Workbench lets you adhere to electrical, physical and timing constraints during the entire design process. In this photo, Allegro's SigNoise (left window) is used to ensure the signal integrity of a trace in the layout (blue and white line at bottom of left window). The graphic display of the noise analysis is shown in the Etchwave tool (right window).

technology and synchronized libraries with EDA tools for design entry, PLD/FPGA design, simulation, physical design, and system-level and circuit board-level analysis, essentially linking back-end analysis tools with front-end tools such as the Verilog-XL simulator.

Setting constraints

A key component of Workbench, the Constraints Editor, lets you set electrical, physical and timing constraints for all the tools in the suite. These constraints are mapped across all the processes in the design, letting you verify that the phys-

example, you can create logical designs without a physical implementation in mind, deferring physical technology decisions until much later in the design process. Using the Allegro CBD option you can apply design-specific and technology-specific constraint definitions of physical designs across multiple packaging and interconnect technologies, such as PCBs, multichip modules (MCMs) and hybrids.

Workbench's closed-loop simulation methodology lets you share timing information throughout the environment. Estimated and actual wire delays, as well as signal-trans-

System Workbench at a glance

- Design suite combines front-end tools for design entry, simulation and analysis with back-end CBD technology and synchronized libraries
- Electrical, physical and timing constraints are mapped across all phases of the design process
- Logical designs can be created while deferring physical technology decisions
- Prance-XL autorouter maps board space into contours to achieve maximum trace density

Cadence Design Systems

2655 Seely Rd
San Jose, CA 95134
(408) 943-1234

Circle 353

NEW AT WESCON/92

5 SHOWS UNDER ONE ROOF

More Products, More
Opportunities to Increase Your
Technical Knowledge, More
Answers to Your Design Questions

INTRODUCING

EC WORLD™
WESCON'S ELECTRONIC COMPONENT SHOW

EC WORLD, the world's largest electronic component event focusing on resistors, relays, capacitors, switches, cable, connectors, enclosures, and associated devices for the design engineer

ITM™
WESCON'S INTERNATIONAL TEST
& MEASUREMENT SHOW

ITM, an expanded world class event for test and measurement, featuring ATE, test and measurement equipment, instrumentation, data acquisition systems, calibration hardware, IEEE I/O devices, and similar products

WESCON'S SEMICONDUCTOR SHOW
icexpo™

IC EXPO, the show for manufacturers of ICs, Semiconductors, ASICs, ROMs, RAMs and similar products

COMPUTOOLS™
WESCON'S COMPUTER TOOLS FOR THE ENGINEERING ENVIRONMENT

COMPU-TOOLS, a computer show for engineers, encompassing hardware and software, EDA tools, workstations, PCs, boards, LANs, CAD/CAM and related products

EPEX™
WESCON'S ELECTRONIC
PRODUCTION EXPO

EPEX, the electronic production expo for worldwide contract manufacturers, surface mount, PCB suppliers, design houses, and products which service this market

You can attend exhibits and technical sessions at any Wescon/92 show with just one admission ticket.

**For Complimentary
Registration**

Call 1-800-877-2668



November 17-19, 1992
Anaheim Convention Center, Anaheim, CA

UNIX-style communications for embedded systems

Embodied Streams from Metasphere brings portable, easy-to-use, UNIX-compatible communication support to embedded systems. The product lets user-level programs, network architectures and higher-level protocols be independent of underlying proto-

AppleTalk, DECnet, Frame Relay, OSI/MAP, SNA, TCP/IP, and X.25/X.29. Communication modules developed for UNIX need only be recompiled and relinked to run in an embedded environment. Embedded Streams conforms to published UNIX V.3 Streams specifications and is written in C to

and VxWorks from Wind River Systems. The size of Embedded Streams' ROMable code is approximately 40,000 bytes.

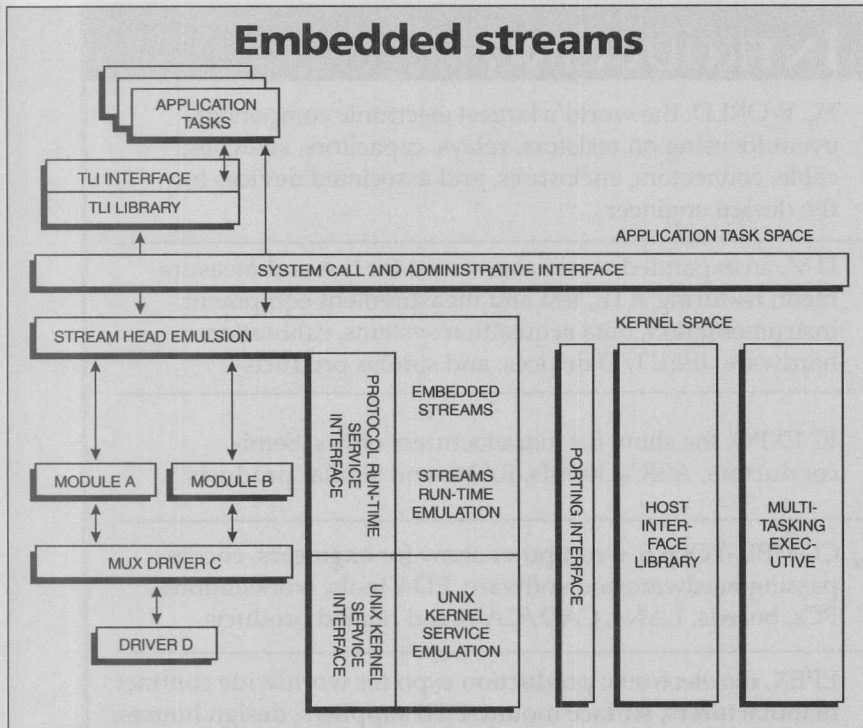
Embedded Streams provides a range of services, including message management and queueing, timer management, buffer management, interrupt processing, flow control, protocol module stacking, asynchronous I/O, message stream multiplexing and demultiplexing, channel polling, and priority message handling.

Also available is an optional Transport Layer Interface (TLI) library with a common program interface conforming to the OSI Layer 4 transport layer. Using the TLI library you can build a variety of protocol stacks, including AppleTalk, ISO-OSI, Netware IPX/SPX, and TCP/IP. Several of these, among them AppleTalk and TCP/IP, can coexist in multiple protocol stacks on the same machine and can be used simultaneously over an Ethernet.

Embedded Streams should make it easier to implement bridges, gateways, hubs, multiplexers, routers, and intelligent network devices.

Product development licenses for the product cost \$8,000. Licenses for the TLI library are \$2,000. Both are shipping now and come with one year of product support.

—Tom Williams



Embedded Streams resides alongside a real-time kernel to which it's adapted via the host interface library. Streams can be constructed by interfacing modules and drivers to the embedded stream head (left). User tasks can access the stream head directly or via a TLI library that offers an OSI Layer 4 programming interface that's common to a number of different communication protocols.

cols and physical communication media.

In UNIX, a stream is a full-duplex processing and data-transfer path between a process in user space and a driver in kernel space. Streams are constructed by interfacing the stream head—the interface between user and kernel space—to various modules such as queues and protocol layers and finally to a driver layer.

Compatibility with UNIX Streams gives embedded systems access to many off-the-shelf communication solutions. Supported protocols include

let user-developed products migrate to new processor families. No licensing from AT&T is required.

Communications executive

Embedded Streams works as a communications executive that resides alongside one of several popular real-time executives. An interface library adapts the Embedded Streams executive to a given real-time executive, or kernel. Executive libraries, available or planned, include Nucleus RTX from Accelerated Technology, pSOS+ from Integrated Systems, VRTX from Ready Systems,

Embedded Streams at a glance

- UNIX-compatible I/O conforms to UNIX V.3 Streams specification
- Interface libraries to popular real-time executives
- 40,000-byte ROMable code size
- Supports protocol families for off-the-shelf communications solutions
- Optional TLI library for easy interface to user applications
- One-year product support included

Metasphere
2560 Ninth St., Ste 121
Berkeley, CA 94710
(510) 540-7015

Circle 354

EMBEDDED COMPUTER CONFERENCE AND EXPOSITION

APRIL 14-16, 1993 SANTA CLARA CONVENTION CENTER, SANTA CLARA, CA

The industry's first and only conference and exposition dedicated to the design and development of OEM embedded computers.

Unlike other industrial conferences and shows which focus on components or technology, ECC concentrates on applications and solutions to embedded computer needs. So, if your system calls for an embedded computer, you'll be able to learn about, and look at, all the alternatives from chip implementations to single-board solutions, to bus-based solutions and even embedded workstations. And most important, you'll learn about all the software, tools and what to do with them.

► System Architecture Issues:

- PCs (80x86, P5) in Embedded Applications
- RISC Architectures (Alpha, HP-PA, MIPS, SPARC, 88K) in Embedded Applications
- Single-board System Implementations
- Custom and Semicustom System Implementations
- VMEbus, Multibus, Futurebus+, STD/STD 32 and Others
- SPARC-Based Single-Board Workstations
- Standard Mezzanine/Daughter Boards
- SBus, TURBOchannel and Other Workstation I/O Buses
- Military Systems and Standards
- And More

► Software Focus On:

- Real-time OS and Kernels
- Real-time DOS
- Windows for Embedded and Real-Time Applications
- Multiprocessing With DOS and Other Real-Time OSs
- High-performance Optimized Compilation
- POSIX and POSIX Compatibility
- Communications Protocols Standards
- And More

► Typical applications

- Medical Instrumentation
- Vehicular Traffic Control
- Automated Vehicles
- Machine Control
- Low-power and Portable Applications
- Visual Inspection
- Military C³I
- Communications
- Peripheral Interface and Control
- Process Control
- Laboratory Automation
- Multimedia
- Networking and More

Plain and simple: If you need an embedded computer in a product you're developing or building, ECC is the one conference and exposition for you. ECC offers the broadest range of solutions to embedded computers ever assembled in one place.

The World-Class technical program includes lectures, application-oriented sessions, workshops and seminars on system architectures, local and subsystem interfaces, I/O interfaces and communications, software and development tools.



EMBEDDED COMPUTER CONFERENCE
& EXPOSITION

*Software
and Hardware
Solutions
for Embedded
Computer
Applications.*

CIRCLE NO. 80

Return the coupon below or call (203) 831-9444, or FAX the coupon to (203) 838-1447.

☐ YES, Send information on exhibiting ☐ YES, Send information on attending

Name _____

Title _____

Company _____

Address _____


City _____ State _____ Zip _____

Phone _____ Fax _____

Mail Coupon to: ECC '93, % BAV Expositions, Inc.

85 Washington Street, South Norwalk, CT 06854

AHDL moves slowly toward standardhood



Peter Denyer, vice-president of marketing for AnaCAD (Fremont, CA), speculates that "analog circuit synthesis is the wave of the future." If he's right, an analog hardware description language (AHDL) is a prerequisite.

There are several organizations attempting to define a usable AHDL: two defense agencies, an IEEE standards group and a number of commercial CAE tool vendors. Each is approaching the problem from a different angle, so where, when and how they'll converge on a standard is an open question, but we hope to cast some light on the issue when Barbara Tuck's panel on AHDL meets at the annual Analog and Mixed-Signal Design Conference in Burlingame, CA later this month.

Presenters at the conference will include people representing:

- SCC-30, the IEEE standards coordinating committee charged with exploring and compiling an AHDL standard;
- Intermetrics (McLean, VA), the primary contractor charged with developing a MIMIC (millimeter and microwave IC) HDL (MHDL), which many believe will be useful for low-frequency analog;
- the VHDL subpar committee charged with developing analog extensions to VHDL;
- Analogy (Beaverton, OR), whose MAST modeling language is considered by some to be a foundation for AHDL development and whose technologists play visible roles in MHDL and SCC-30 working groups; and
- Cadence Design Systems (San Jose, CA), a SCC-30 participant with its own view of the direction AHDL development should take.

Tuck's panel promises only a snapshot of what these groups are working on. Development efforts may take a different turn in the future, especially since the U.S. Air Force has indicated a willingness to sponsor talks on AHDL beginning in 1993.

The Air Force's participation is significant because it was a DOD initiative that drove development of VHDL

as both a government procurement standard and a commercial specification. By ensuring that system descriptions were in the same computer-readable language, the government could solicit bids from a wide variety of contractors without fear of ambiguity about what the proposed system would do.

Dr. Joel Schoen, chairman of SCC-30 and simulations director at Mitre Corporation (Bedford, MA), has said that development of an analog language will be driven by the DARPA-financed effort to develop an MHDL. The MHDL contract, administered by the U.S. Army LABCOR Electronic Technology and Devices Lab (Ft. Monmouth, NJ), aims to develop a language for microwave systems, but since the MHDL working group is staffed by analog gurus such as John Paulos of North Carolina State and Ernst Christen, engineering manager of simulation at Analogy, Intermetrics (the language developer) is making an extended effort to build something that will be useful for analog designers. With funding and a schedule, Schoen suggests, MHDL development will go faster than the strictly voluntary efforts that characterize SCC-30 and other IEEE committees.

■ Military efforts redundant?

It's unclear whether the proposed Air Force contract on AHDL, administered by Rome Laboratories (Griffiss Air Force Base, NY), will be redundant with the Army's MHDL contract or will take an entirely different approach. Air Force contract administrators maintain they have no intention of competing with the Army or any other group on AHDL. But, many of the same organizations participating in the Army's MHDL working group have already asked to be included in Air Force-sponsored discussions, among them Analogy, Cadence Design Systems, CLSI, and Intermetrics. Whatever the outcome of the Air Force-sponsored talks, it must be presented to the IEEE as a useful and desirable standard. Consequently, the Air Force-sponsored discussions may serve only to ratify the MHDL design language, or could come up with analog extensions to VHDL, or may originate something entirely different.

In many ways, the direction of the Air Force talks

PALMTOP & CELL PHONE SUPPLY STARTS FROM 1.8V!

Surface Mount IC Draws Only 1.6mA, Delivers 5V at 1A

The MAX741U's unique combination of low start-up voltage and low quiescent current make it unbeatable in battery-operated and portable power supply designs that require high output currents from low input voltages. Use the MAX741U in applications such as:

- **Disk Drives:** Step-Up to 5V at 1A in 3V-Only Systems
- **Palmtops:** Deliver 5V at 500mA from 3V Batteries
- **Cellular Phones:** Use SYNC Pin to Deliver Clean Power to Transmitters


EVALUATION KIT SUPPLIES 5V AT 1A FROM A 2.7V INPUT

1.6mA Quiescent Current
(50µA in Shutdown)

**145kHz PWM
Current-Mode Control**
(Easy to filter fixed-frequency
ripple & excellent transient
response)

External Oscillator SYNC Pin
(For noise-sensitive cellular phone
& data acquisition applications)

**Single-Ended,
Complementary,
& Push-Pull Outputs**
(For single-inductor, synchronous
rectification, or transformer
applications)



**Undervoltage Lockout
& Soft-Start Protection**

A complete pre-assembled surface-mount evaluation kit demonstrates disk-drive and palmtop applications: stepping up to 5V at 1A from a 2.7V input—order the MAX741UEVKIT-SO kit. For 5V, 1.5A output step-down applications, order the MAX741DEVKIT-SO kit. \$30.00 each.

Low Supply Current MAX741 Also Steps-Down and Inverts

| | Function | Output Voltage (Pin-Selectable) | Quiescent Supply Current | Features | Price* |
|---------|-----------|------------------------------------|-----------------------------|--|--------|
| MAX741U | Step-Up | 5V, 12V, 15V, Adj. | 1.6mA | Evaluation kit, 1.8V start-up voltage | \$3.25 |
| MAX741D | Step-Down | 5V, Adj. | 2.8mA | Evaluation kit, 90% efficiency, 5V at 3A from 6V input | \$3.25 |
| MAX741N | Inverting | -5V, -12V, -15V, Adj. | 2.2mA | 1A output current, power ECL logic | \$3.25 |



FREE Power Supply Design Guide – Sent Within 24 Hrs!

Includes: Data Sheets and Cards for Free Samples

CALL TOLL FREE 1-800-998-8800

For a Design Guide or Free Sample

Master Card® and Visa® are accepted for Evaluation Kits or small quantity orders.

MAXIM

Maxim Integrated Products, 120 San Gabriel, Sunnyvale, CA 94086, (408) 737-7600, FAX (408) 737-7194.

Distributed by Arrow, Bell/Graham, Elmo, Hall-Mark, Nu Horizons, and Pioneer. Authorized Maxim Representatives: **AL**, M2i Montgomery Marketing, Inc.; **AZ**, Techni Source Inc.; **CA**, Mesa, Pro Associates, Inc.; Centaur Corporation; **CO**, Component Sales; **CT**, NRG Limited; **DE**, TAI Corporation; **FL**, Sales Engineering Concepts; **GA**, M2i Montgomery Marketing, Inc.; **ID**, E.S. Chase; **IL**, Heartland Technical Marketing Inc.; **IN**, Technology Marketing Group; **IA**, JR Sales Engineering, Inc.; **KS**, Delltron; **LA**, BP Sales; **MD**, Micro-Comp, Inc.; **MA**, Comp Rep Associates; **MN**, Mel Foster Technical Sales, Inc.; **MS**, M2i Montgomery Marketing, Inc.; **MO**, Delltron; **MT**, E.S. Chase; **NE**, Delltron; **NV** (Reno, Tahoe area only) Pro Associates, Inc.; **NH**, Comp Rep Associates; **NJ**, Emtec Sales, Inc.; TAI Corporation; **NM**, Techni Source Inc.; **NY**, Parallax, Reagan/Compar; **NC**, M2i Montgomery Marketing, Inc.; **OH**, Lyons Corporation; **OK**, BP Sales; **OR**, E.S. Chase; **PA** (Pittsburgh area) Lyons Corporation, (Philadelphia area) TAI Corporation; **SC**, M2i Montgomery Marketing, Inc.; **TN**, M2i Montgomery Marketing, Inc.; **TX**, BP Sales; **UT**, Utah Component Sales, Inc.; **VA**, Micro-Comp, Inc.; **WA**, E.S. Chase.

Distributed in Canada by Arrow and Future. Authorized Maxim Representative in Canada: Tech Trek.

*MAX741U, MAX741D, or MAX741N. 10,000up, US recommended resale.

Maxim is a registered trademark of Maxim Integrated Products. © 1992 Maxim Integrated Products.

MIXED-SIGNAL DESIGN

will be contingent on the results of the Army's efforts to develop an MHDL useful for analog design work. The Army maintains that it's evolving a general-purpose language that should be useful for low-frequency analog specification work. Casual observers wonder, however, if there isn't a crossroads between the needs of microwave system builders and those of low-frequency analog designers—that is, those who work below 100 MHz. The MHDL working group includes radar system builders such as Raytheon and Mitre, radio frequency and microwave tool vendors such as EEsof and Compact, and chairperson Robert Bierig, a microwave consultant for B&B Technologies (Hopkinton, MA).

■ Microwave and analog

Microwave design constructs don't always converge with those of low-frequency analog. Because almost everything that happens above 100 MHz must be viewed as some sort of transmission line effect, microwave designers insist on references to physical geometries in their designs. Analog designers, on the other hand, utilize electrical constructs—KCL (Kirchoff's Current Law) and KVL (Kirchoff's Voltage Law)—to predict the currents and voltages about a node; these constructs aren't particularly useful for microwave systems design. Consequently, analog spokesmen Paulos and Christen have had to lobby hard to keep conservation laws such as KCL and KVL strong in the MHDL Requirements Document that Intermetrics is using for its guideline.

Considering the potential schism between the needs of microwave and low-frequency analog designers, Intermetrics has done a masterful job of pulling together high-level language constructs that may be useful to both groups. Much of what follows should be regarded as work in progress, which may have changed radically by the time the Analog and Mixed-Signal Conference convenes.

Senior computer scientist Dave Barton of Intermetrics is working on a signal algebra, a method of describing analog and microwave behavior in terms of signals. The MHDL, suggests Barton, must provide a shorthand for describing a signal as well as the hardware that produces it. A language based on transfer functions, he reasons, brings us close to where we need to be.

In its first attempts at constructing an MHDL, Intermetrics has utilized concepts from a new functional language called Haskell (named after the logician Haskell B. Curry and developed by 15 contributors from Los Alamos, MIT, the University of Glasgow, and Yale). The advantage of Haskell is that it describes higher-order signal processing functions with user-definable data types. It lets both stream and continuous-data (or continuation) models be processed with extensible operators.

One of the first Haskell implementations was put together at Yale. It runs on the Sun-4 SPARCstation, and is compatible with Emacs, the signal-processing program originated at Leuven University in Belgium. (Both Mentor Graphics' DSP Station soft-

ware and Star Semiconductor's DSP development tools can trace their roots to Emacs.) As embodied in MHDL, Haskell concepts encourage you to define classes of signals, components and connections.

■ Developing a language

In early versions of MHDL, language syntax is considered to be relatively unimportant. Barton, for example, has presented versions to MHDL working group members and asked that they be tried on for size. The language development process, then, is largely iterative—Barton presents guidelines for MHDL statements, members go offline to try them out and the testers return monthly to report their results. If working group members report good results, Intermetrics feels it's on the right track. If members report problems, then Barton identifies the problem area and attempts a modification.

What this process tells us is that completion of a full MHDL is still a long way off. Many assumptions about the utility of Haskell in particular and a signal-processing algebra in general may change in the interim. But it's clear that the Army and its contractor are indeed making an earnest effort to satisfy the needs of low-frequency designers.

In his first attempt to construct a motor control circuit in MHDL, Analogy's Christen complained that Haskell is a language designed by a committee of language theorists for other language theorists. He wrote recently, "I spent several hours with [the Haskell tutorial] trying to learn about Haskell, but at the end was more confused than enlightened." About MHDL he has written: "Reading through the MHDL Preliminary Language Reference Manual leaves the feeling that I had learnt about many pieces of a puzzle, but I had not been given any instructions on how these pieces fit together."

Indeed, in rereading the minutes of the June MHDL working meeting, I was reminded of Cadence product marketing manager Eric Filseth's observation that analog designers are not software jocks and would resist learning a new language. Knowledgeable industry participants invited to sit in on the June working group meeting and assess MHDL for the first time expressed a sense of being overwhelmed by what they were seeing.

Progress on MHDL and AHDL, and feelings about such progress, will undoubtedly change by the time Tuck's panel convenes. Remember that the effort of building a standard is both difficult and painstaking, involving many hours of detailed work. The MHDL working group, the AHDL working group and IEEE committee members are attempting to construct something that will be useful for all of us. So when you meet these people at the Analog and Mixed-Signal Design Conference, please shake their hands and give them your applause.

Stephan Ohr is president of Indian Forest Research and publisher of the monthly newsletter, Mixed Signals.

COMPUTER DESIGN is now organizing the technical programs for three major industry events with many opportunities for your company or your clients to participate. Each of these world-class technical programs will have a mix of tutorials, individual lectures, targeted multi-presentation sessions, workshops, and panel discussions.

If you're a supplier of hardware, software or development tools, each of these technical programs is an ideal opportunity to put your technical expertise on display and reach a dedicated group of designers and decision makers. Or if you're an OEM or systems house that's developed an innovative solution or application of RISC, fuzzy logic or other embedded computer technology, these conferences are your chance to share some of what you've learned.



RISC '93
MARCH 16-18, 1993
HYATT REGENCY, SAN FRANCISCO AIRPORT

Under the guidance of senior editor Don Tuite, the RISC '93 program will focus on advances in RISC processor architectures, processor/memory interfaces, multiprocessing, RISC in real-time, peripheral chip sets, operating systems and kernels, compilers, development tools, and benchmarking RISC performance. This three-day, total-immersion conference will give designers and design managers an understanding of all the major RISC architectures and the design trade-offs in their use.



EMBEDDED COMPUTER CONFERENCE (ECC)
APRIL 14-16, 1993
SANTA CLARA CONVENTION CENTER

The ECC technical program is dedicated to the major architectural aspects of embedded computer systems including processor and memory architectures; MCM and module interfaces; real-time and multiprocessing software solutions; interconnect structures and standards, including buses; and I/O interfaces and standards. The technical program chairman is senior editor Warren Andrews.



FUZZY LOGIC '93
JULY 20-22, 1993
HYATT REGENCY, SAN FRANCISCO AIRPORT

The technical program for Fuzzy Logic '93, directed by senior editor Tom Williams, brings fuzzy logic out of the world of academia and into the real world of available products, tools and applications. The conference will focus on fuzzy logic processors, coprocessors and board-level implementations; fuzzy logic architectures and interfacing; development tools; and fuzzy logic in specific applications such as image processing and pattern recognition, insurance and market analysis, health care and medical diagnostics, automotive systems, and weapons guidance and control.

Full details about each of these technical conferences and how you can participate, is contained in the appropriate Call for Papers. To receive the Call for Papers for a specific conference, contact:

PATTI KENNEY, TECHNICAL PROGRAMS COORDINATOR, AT (508) 392-2124



A PennWell Publication

CIRCLE NO. 82

8051 IN-CIRCUIT EMULATOR


\$995⁰⁰

- 8031/32/51/52 support
- 64K Memory/64K breakpoints
- Source Level C and Assy debugging
- Trace Buffer ■ Programmers
- Money Back Guarantee

VAIL SILICON TOOLS
305 570 5580

CIRCLE NO. 166

☆☆☆☆
"The Best 8051 Emulator"



5 ft. cable

8051

SEE EEM 91/92
Pages D 1300-1303

PC based emulators for the 8051 family

8031, 8032, 8051, 8052, 80C152/154/321/451/452/51FA/51GB/51H/51J/51K/51L/51M/51N/51O/51P/51Q/51R/51S/51T/51U/51V/51W/51X/51Y/51Z/52, 8751, 8752, DS5000 + CMOS... more.

- PC plug-in boards or RS-232 box.
- Up to 33 MHz real-time emulation.
- Full Source-level Debugger w/complete C-variable support.
- 64 bit wide, 256K deep trace, with time stamp.
- Bond-out/hooks pods for 8051, 83C552, 83C451, 83C652, 83C751, 8XC51/80C517, 83C752, 8XC51FA/FB/FC, and more.

Prices: 32K Emulator 8031 \$1790; 4K Trace \$1495* (*US only)

CALL OR WRITE FOR FREE DEMO DISK!

NOHAU
CORPORATION

51 E. Campbell Avenue
Campbell, CA 95008
FAX (408) 378-7869
(408) 866-1820

Call 408-378-2912
Nohau's 24-hour
information center to
receive info via your FAX

CIRCLE NO. 167

REMOVE HARDWARE LOCKS

Protect your investment!

Don't wait for your lock to get lost, stolen or fail and render your program useless.

Our **SAFEKEY'S** are guaranteed to unlock the following:

PADS 2000 MAXROUTE TOPAS
PADS PCB MICROSTATION RIO
SMARTCAM 3D STUDIO & MORE

From \$399 plus S&H
(609) 390-2799

24 HR. FAX HOTLINE
(609) 390-3750




IMAGINE THAT

1 Evergreen Dr., Seaville, NJ 08230
TM-Trademarks property of respective owners.

CIRCLE NO. 168

NEW! TIMING DIAGRAMS



dV/dt® Timing Analysis Tools for your PC or MAC

Now, draw, analyze, modify and document digital circuit timing diagrams and designs.

- Fast
- Easy
- Accurate

Only **\$495**

(includes Test Vector Generator)

Call 619-292-1900 for FREE DEMO

engineerium

(computer product development)

7968-104 Arjons Dr., San Diego, CA 92126 • 619-292-6236 FAX

CIRCLE NO. 169



357 MHz 14-Pin Clock Oscillator

Connor-Winfield Corp. announces literature on a new line of high frequency ECL Logic clock oscillators. The ECLB Series oscillator covers the high frequency range of 8 MHz-360 MHz. AVAILABLE NOW IN A 14-PIN DIP. Literature for a double DIP version with frequencies to 500 MHz is also available. Frequencies to 90 MHz are now available with Voltage Control function (model EV53 series). Contact Barney Ill for details.

Additional specifications listed below:

Model : ECLB
Package : 14-pin DIP
Frequency: 8 MHz to 360 MHz
Supply : -5.2 or -4.5 Vdc

Connor-Winfield Corporation

1865 Selmar Road
Aurora, IL 60505
PH: (708) 851-4722 FAX: (708) 851-5040

CIRCLE NO. 170

SYSTEM SHOWCASE


Reach over 100,000
qualified engineers and
engineering managers
with
SYSTEM SHOWCASE
advertising.

Rates start at \$940

For more information
call Sue Shorrock at
800-225-0556 or
508-392-2185

CIRCLE NO. 171

Instant Microcontroller



Instant C Programming

Don't use a microprocessor, use a **SmartBlock™** microcontroller module to build your custom controller. Our \$195 interactive **Dynamic C™** development system makes programming easy. 3.5 x 2.5 inch module includes microprocessor, memory, time/date clock, eeprom, watchdog, serial ports and more. As low as \$59 in quantity. The efficiency of a custom design without the headaches.

Z-World Engineering

1724 Picasso Ave., Davis, CA 95616 USA

Tel: (916) 757-3737

Regular Fax: (916) 753-5141

Automatic Fax: (916) 753-0618

(Call from your fax, request data sheet #14.)

CIRCLE NO. 172

High Performance DAC-Per-Pin™ Device Programmers



Programs virtually every device available today with a certified programmer from the leading U.S. supplier of hardware/software device tools. The Allpro starts at \$1295 and is fully upgradeable to 88 pins.

We Have the Best Products on The Market.

For more information,

call: 1-800-331-7766

or (305) 428-6868.

**LOGICAL
DEVICES, INC.**

CIRCLE NO. 173

State Machine Design

For Complex & High Density PLDs



The most powerful PLD/FPGA CAE
design software from \$495.00

CUPL™ 4.2

1-800-331-7766


305-428-6868

**LOGICAL
DEVICES, INC.**

CIRCLE NO. 174

EPROM PROGRAMMER

\$750



- Quick pulse pgms. eight 1 Mbit EPROMs in 40 sec.
- Stand-alone or PC-driven
- RS-232, parallel in & out ports
- Binary, Intel hex, & Motorola S formats
- 100 user-definable macros
- Information, call (916) 924-8037
- 1 Megabit of DRAM
- Made in U.S.A.
- A9 Identifier
- 2 year warranty
- Single pgmr. \$550

NEEDHAM'S ELECTRONICS

4539 Orange Grove Ave. • Sacramento, CA 95841
(M-F, 8-5 PST)



CIRCLE NO. 175

GMS V49

FEATURE
COMPARABLE
MME 147
PERFORMANCE
UPGRADE

68040-Based, Real-Time,
Multiprocessing SBC

I/O Features:

- Ethernet™
- SCSI
- 8 Channels of SIO

Dual-Ported Memory

- 2, 4 or 8MB of DRAM (2-way interleaved, w/ burst)
- Up to 1MB of b-b SRAM
- Up to 4MB of EPROM

The V49 also features the VIC068 VLSI VMEbus interface and a RTC

A 68030 version of this SBC is also available




GENERAL MICRO SYSTEMS
8358 Maple Place
Rancho Cucamonga, CA 91730 USA

(714) 980-4863 • FAX 987-4863

CIRCLE NO. 176

Tiny Embedded 386SX

-40° to +85° C



Model 5025

- ▶ 386SX, 25 MHz
- ▶ 4.5" x 4.9"
- ▶ 3 solid-state disks
- ▶ Low power CMOS
- ▶ Rugged, industrial quality
- ▶ Stand-alone or bus operation
- ▶ Flash memory programmer
- ▶ Coprocessor socket
- ▶ COM1, COM2, LPT1 ports
- ▶ 4 MB DRAM
- ▶ "Instant DOS" 3.31 in ROM
- ▶ Autotest operation
- ▶ SV only operation
- ▶ Keyboard/speaker port
- ▶ Watchdog timer
- ▶ Calendar/clock

THE MICROCOMP COMPANY

OCTAGON SYSTEMS
Tel: 303-430-1500, Fax: 303-426-8126

Call or Fax for full-line catalog

CIRCLE NO. 177

FINALLY

One tool to satisfy all your
firmware development needs

PROMICE is a universal system.

- Develops code for any microprocessor
- Complete, real time, source level debugging
- Host software for DOS, Unix, Mac, VMS
- Non intrusive on your target system
- Simply plugs into any ROM socket

PROMICE also supports Turbo Debugger, C_thru_ROM, FreeForm, GDB, and more.


PROMICE

...the affordable solution.



TEL: (614) 899-7878
West: (415) 750-0219
FAX: (614) 899-7888

CIRCLE NO. 178



IT CAN BE THIS EASY!

OrCAD's innovative ESP Framework can make creating a netlist from your schematic as easy as pressing an on-screen "button."


Call (503) 690-9881 for a FREE demonstration disk and information about the entire OrCAD product line: schematic design, simulation, pc board layout, programmable logic IC design and more.

CIRCLE NO. 179

386/486 Bolt-in Computers™ for embedded applications

High performance CPUs for aggressive applications. From full blown VGA flat panel, 8MB RAM, 510 MB IDE systems. To number crunching high throughput ROM-based systems.

- 16-25 MHz 386SX/486SLC
- 8MB RAM, 1.8 MB ROM
- Solid state or IDE disk for DOS 5.0, Windows 3.1*
- Multitasking, multidrop comm, PID control, multi-CPU software
- Optional STD BUS expansion



From \$995,
quantity 1, great OEM discounts.

MICRO/SYS

3447 Ocean View Blvd., Glendale, CA 91208
(818) 244-4600 FAX (818) 244-4246

CIRCLE NO. 180


TARGET COMPUTER DESIGN SUBSCRIBERS BY MAIL

COMPUTER DESIGN subscriber list available for rent. 100,000 direct mail responsive engineers and engineering managers. Key decision makers, by name, at business addresses. Target by job function, company type, design management, product design, purchase influence or geographic areas. Proven winner for books, subscriptions, technical reports, seminars, conferences, tools, components, catalogs, hardware, software, testing instruments and many other offers.

For more information call
Deanna Rebro at PennWell Lists,
Advanced Technology Group:

800-982-4669 or 918-831-9551

CIRCLE NO. 181



SBX ANALOG I/O GET A COMPETITIVE EDGE!

Density - up to 16 A/D inputs & 8 D/A outputs on one single width card. Speed - throughput rates up to 59 kHz. Intelligence - many preprogrammed modes. Input filters, prog. gain amps, sample-holds, FIFO I/O buffers & many other features too numerous to mention here.

ROBOTROL CORP

925 W. San Martin Ave, PO Box 990
San Martin, CA 95046
(408) 683-2000

CIRCLE NO. 182

HIGH PERFORMANCE GRAPHICS CARDS

- 32/40 MHz TMS34020 GSP
- Optional FPP w/512K local SRAM
- For ISA Bus & STD 32® systems
- 1 MB DRAM, 1 MB VRAM standard
- On board memory expandable to:
 - 4 MB DRAM
 - 4 MB VRAM
 - 512K EPROM
- 640x480 thru 1600x1200 pixels
- 24-bit True Color mode
- 256 colors from 16.7 million
- Free TIGA® & Windows® drivers


\$1,495 up

ZWICK SYSTEMS INC.

Tel 613-726-1377, Fax 613-726-1902

CIRCLE NO. 183

Half size ECL Clock Oscillator



Connor-Winfield Corp. introduces its new 8-pin DIP ECL clock oscillators. The E500 Series is available in frequencies from 24 MHz - 180 MHz. Options available include industrial temp range of -40C to +85C, complimentary output, enable/disable function, and surface mount header. Frequency stabilities are available to +/-25 ppm over temperature.

These models are available with supply voltage of either -5.2 V dc (10K equivalent) -4.5 Vdc (100k equivalent) or +5Vdc. Prototype quantity pricing at 120 MHz is \$43.90 each. Delivery is stock to 7 weeks A.R.O. Contact Barney III at: **Connor-Winfield Corporation** 1865 Selmarten Road, Aurora, IL 60505
PH: (708) 851-4722 FAX: (708) 851-5040

CIRCLE NO. 184

CapFast Schematic Capture and Interface Tools

PRICE STARTS AT \$595.00

Available on both DOS (PC) and UNIX (SPARC, DEC, HP)

CapFast is an advanced hierarchical schematic design and interface tools for PCB, PLD/FPGA, as well as ASIC designs.

The optional EDIF 200 tools allow the designer to translate CapFast schematics to and from other EDA systems, such as Mentor Graphics and Cadence. For further information, please call (503) 645-0313.



Phase Three Logic

CIRCLE NO. 185

QUICK-TURN PC BOARDS

INSTANT QUOTES

1-800-234-1556

(After Hours (602) 377-7319)

Single/Double Sided/Multi-Layer
Gerber Input/Floppy or Modem



MIL
Approved
U.L. Listed



AVANTI CIRCUITS, INC.



17650 North 25th Avenue
Phoenix, Arizona 85023

(602) 863-7729 & 234-1556

Fax (602) 375-1909

BBS/Modem (602) 234-1737

CIRCLE NO. 186

EMPLOYMENT OPPORTUNITIES, OUTPLACEMENT AND POSITIONS WANTED

Recruitment Advertising

Sue Shorrock, (508) 392-2185

CAREER CONNECTIONS

Closing Dates for upcoming magazine issues

| Issue | Closing |
|-------|---------|
| NOV | OCT 4 |
| DEC | NOV 4 |

POSITION WANTED ADS

Free 1" ad to subscribers seeking full-time employment. Just include 50 words of copy and your subscription label. We'll run your ad in 2 consecutive issues. Available to non-subscribers or consultants/companies at \$125 per column inch. Mail your position wanted ad to: **COMPUTER DESIGN Positions Wanted**, One Technology Park Drive, PO Box 990, Westford, MA 01886



POSITIONS WANTED

GREATER BOSTON, MA AREA: Interested in adding an experienced COMPONENT/PRODUCT ASSURANCE or MANUFACTURING ENGINEER to your staff? Look no further! Over ten years experience: BS Electronic Engineering Technology; Supervised staff of 8; Military/commercial manufacturing; Call George J. Alexandropoulos (508) 692-7301. I'll FAX resume.

SEASONED PR/ADVERTISING/MARCOM MANAGER, 26 years experience with many awards, seeks sr. level corporate or agency position. Heavy magazine PR experience, plus advertising, collateral, trade shows, videos, newsletter.

Former radio/TV reporter. Experienced at planning, managing, executing. Exceptionally good at developing and managing individual contributors. Live Northeast, will relocate. Call Jeff, (508) 746-7898.

TECHNICAL WRITER. Eight years publications experience. Documentation and promotions including copy and scriptwriting. Knowledge of popular PC packages. Clients include CAD, robotics, software developers, design. Offer free diagnosis of your publications needs. Samples. Call Jay at (617) 523-4217.

SENIOR LEVEL MULTIMEDIA SYSTEMS MARKET ANALYST, STRATEGIC PLANNER AND PRODUCT DEVELOPER. Member of ANSI and ISO Experts Groups on: digital Image, Motion Video, Audio compression, and ISO Electronic Still Pictures Imaging. Experience developing MPEG Video chip, DSP-based MPEG-Audio software solution, Laser Disc-based Multimedia Applications. Call Alfred Ricconi (214) 644-8875.

Seeking Regional Sales Manager position. Self motivated individual with extensive and successful experience selling technical products and services to industrial customers in the southwest and Rocky Mountain regions. Products sold include both hardware and software for industrial control applications. Additional experience as software engineer. Contact: Nick Baker, 2015 Martingale Dr, Norman, OK 73072. (405) 364-4628.

MICROPROCESSOR EXPERT, Hardware/Software, Optimized Assembly, C, 680x0, Emulators, Multiprocessing, Device Driver's, PCs, Graphics, VME, Realtime, Instrumentation, 16 Bit Analog, Signal Conditioning, Semiconductor Testing, DSP, Industrial Control, Servo Systems, Prototypes, Project Management, Total Life Cycle. John Odendahl, 7801 Shoal Creek #112, Austin, TX, 78757, (512) 458-8856.

PROJECT MANAGER: Highly versatile professional with over 15 years of broad-based experience in Project Management, Economic Analysis and Engineering gained with blue chip firms (GTE, PacBell, FHP, etc.). Technically oriented, MBA, BS, and BA degrees with strong functional skills embracing all areas of research, applications engineering and business analysis. Call Bob Mazzola, (714) 841-8978.

Fault tolerant computing, digital, analog, software hands-on design and analysis. Reliability, cost and performance modeling. Failure Modes and Effects, worst-case analyses. Multi-tasking software and hard real-time systems design. SOA system methodologies. Don Uhrich, DEPENDABLE SYSTEMS, 7088 Olde Stage Road, Boulder, CO 80302, (303) 440-4696.

SENIOR ENGINEERING MANAGER/ENGINEERING DIRECTOR: Over 20 years of successful and professional experience in major industrial electronic companies. Expertise in planning conceptual electronic design, with emphasis on managing, directing and developing people for project and product success. BSEE, MSEE and MBA. Resume on request. Bruce Carlton, 2045 E. Knox Rd, Tempe, AZ 85284, (602) 831-9699.

Product Development: Hardware, Firmware, Software. Developed and owns *complete rights to market-ready* Data Acquisition/Voice/Phone/Instrumentation products. Experience in Embedded Controllers, Assembly language, PC interface and programming. Schematics Design, PCB layout, prototype development, & troubleshooting. Brochure, advertising and manual write-up. Excellent communications skills. Contact Nidal at (719) 260-6834.

PROGRAM/PROJECT MANAGER. Over 15 years experience in the design and development of large scale data acquisition and control systems. Managed and performed real time software & hardware systems design. Background in instrumentation, control, data acquisition, signal conditioning, transducers etc. Moorer, 1000 Mira Vista Dr, Huntsville, AL 35802 (205) 881-8444.

ENGINEERING/APPLICATION/CUSTOMER ENGINEERING. BSEE with 20 years of hands-on management of design from concept to delivery. Customer/vendor contact, Analog/Digital/Packaging/ and use of CAD. Peripherals/small system exp. Kayvon Kiamanesh, 21811 Satcoy St # 21, Canoga Park, CA 91304.

Q.A. Manager seeks position with a company desiring to improve its performance, i.e., quality & productivity. I have excellent related experience that includes quality program audits, proposals, plans, procedures, TQM, SPC, SQA, QE, MRB, Config. Mgt. and ISO programs for aircraft & aerospace products & services. Please call (818) 365-2200 if no answer please leave message.

PRODUCT DESIGN MANAGER/PARTNER. Does your company have the need for a custom embedded Basic like interpreter? Also do digital/analog/custom software development for both PCs and real time systems. Southern California preferred. Larry Dersh (619) 670-3817, 10258 Fairhill Dr, Spring Valley, CA 91977.

CHALLENGES

E-Systems ECI Division is in need of engineers with military satellite communication experience (ground, manpack/manportable, airborne, or spaceborne); a BSEE/BSCE; and at least 2 years' experience in one of the following areas:

SOFTWARE

- ADA, C
- 1750A, of 68020 Microprocessors
- VAX, VMS, Sun UNIX
- Real-time, Embedded Microprocessor
- DOD-STD-2167A, CASE Tool

DIGITAL HARDWARE

- ACTEL FPGAs
- Microprocessor based systems
- 1553 bus interface

DIGITAL SIGNAL PROCESSING

- Discrete Fourier Transforms
- Control Loops
- PSK Demodulation

EMBEDDED CRYPTO

- Security Fault Analysis
- TEMPEST, Red/Black Isolation
- Related interface hardware

RF AND MICROWAVE

- Synthesizer Design, Direct Digital
- Power amp and filter design
- MMIC design

ANTENNA DESIGN

- Parabolic Antenna Design
- Gimbal, Positioner
- 10 TO 60 GHz

SYSTEMS

- BSEE/MSEE, minimum 4 years' experience
- Strong communication background
- Requirements Analysis, Functional Analysis
- System Synthesis, System Analysis
- RF Link Budget Analysis
- System Integration/Test
- Customer Interface

Background: Hardware, Software Architecture Cryptographic; BIT/BITE; Antenna Pointing, Tracking, and Platform Stabilization; MIL-STD-1582; SI-1135, SI-2035, LL1005.

MILSTAR or other military satellite design experience a plus.

E-Systems offers very competitive salaries and an excellent benefits package which includes an Employee Stock Ownership Plan, 401(k), and major medical and dental insurance. Qualified candidates should forward a resume and salary history to: Manager of Staffing, E-Systems, Inc., ECI Division, Post Office Box 12248, St. Petersburg, Florida 33733-2248.




E-SYSTEMS
The science of systems.

U.S. Citizenship Required.
An Equal Opportunity Employer, M/F/D/V.

40 MB/SEC DUAL-PORT VME/VSF HIGH - DENSITY MEMORY

from
Chrislin Industries, Inc.
Your
TOP QUALITY Memory Supplier
for
OVER 16 YEARS



CI-VME40

- ☐ **FAST ACCESS TIMES**
20ns READ/WRITE BLOCK CYCLE
- ☐ **FAST CYCLE TIMES**
83/62NS READ/WRITE BLOCK CYCLE
- ☐ **VERSATILE CONFIGURATIONS**
4, 8, 16, 32 or 64MB in one 6U slot
- ☐ **VERSATILE ADDRESSING**
Addressable in 24 or 32 bit through
4 Gigabytes
Memory start and end addresses
selectable on 256kb boundaries
VME/VSF configured independently
- ☐ **RELIABLE**
VME Revision C.1 Compatibility
- ☐ **LIFETIME WARRANTY**

**OTHER QUALITY VMEbus
MEMORIES AVAILABLE ARE:**

THE CI-VMemory
LOW-COST VMEbus BYTE
PARITY MEMORY with 4, 8, or
16MB in one 6U VMEbus slot

THE CI-VSB EDC
Dual-port VMEbus/VSF memory
with Error Detection and Correction,
single-bit error detection/correction,
double-bit error detection, 4MB up
to 64MB in one VMEbus/VSF slot



Chrislin Industries, Inc.

31312 Via Colinas, Suite 108
Westlake Village, CA 91362

TEL: (818) 991-2254 FAX: (818) 991-3490

CALL TOLL FREE: (800) 468-0736

CIRCLE NO. 102

ADVERTISERS INDEX

| ADVERTISER | PAGE NO. | CIRCLE NO. | **IDA NO. | ADVERTISER | PAGE NO. | CIRCLE NO. | **IDA NO. |
|---|-------------|---------------|--------------|--------------------------------------|-------------|---------------|--------------|
| A | | | | F | | | |
| Advanced Micro Devices | 14-15 | 9 | | John Fluke Mfg. Co., Inc. | 25 | 56 | |
| Advanced Micro Devices | 6-7 | 5 | | Force Computers, Inc. | 8-9 | 6 | |
| Alacron, Inc. | 126 | 69 | | G | | | |
| Analog & Mixed Signal Conference | 50 | 35 | | General Micro Systems | 147 | 176 | |
| Analog & Mixed Signal Conference | 128 | 71 | | Grammar Engine, Inc. | 147 | 178 | |
| *Analog & Mixed Signal Conference | 151 | 85 | | Greenspring Computers | 133 | 75 | |
| AMP, Inc. | 93-96 | — | | Greenspring Computers | 135 | 76 | |
| AP Labs | 132 | 73 | | Greenspring Computers | 137 | 78 | |
| Applied Microsystems | 112A | — | | H | | | |
| AT&T Micro- electronics | 120-121 | 66 | | Heurikon Corporation | 18 | 11 | |
| AT&T Microelectronics | 36-37 | 25 | | Heurikon Corporation | 32 | 20 | |
| AT&T Microelectronics | 49 | 34 | | Heurikon Corporation | 72 | 49 | |
| Augat, Inc. | 63 | 44 | | *Hitachi America, Ltd. | 86 | 57 | 1034 |
| Avanti Circuits, Inc. | 148 | 186 | | Huntsville Microsystems | 4 | 4 | |
| B | | | | I | | | |
| Brooktree Corporation | 53 | — | | IEE Books | 131 | 72 | |
| Brooktree Corporation | 75 | 50 | 1026 | Imagine That | 146 | 168 | |
| C | | | | Integrated Device Technology | 11 | 7 | |
| Chrislin Industries, Inc. | 150 | 84 | | Integrated Systems, Inc. | 71 | 48 | |
| Ciprico | CV2 | 1 | | †Intel Corporation | 151 | 85 | |
| †Computer Design | 33 | 22 | | Intel Corporation | 19 | — | |
| †Computer Design | 37 | 26 | | Intel Corporation | 22-23 | 14 | |
| Computer Design | 105 | 60 | | Invitational Computer Conference | 136 | 77 | |
| Computer Design | 152 | 87 | | Ironics, Inc. | 31 | 19 | |
| Computer Design's Technical Programs | 145 | 82 | | K | | | |
| Compass Design Automation | 106 | 61 | | KADAK | 18 | 12 | |
| †Conference Proceedings | 120-121 | 67 | | L | | | |
| Connor-Winfield Corp. | 146 | 170 | | Logical Devices, Inc. | 146 | 173 | |
| Connor-Winfield Corp. | 148 | 184 | | Logical Devices, Inc. | 146 | 174 | |
| CSPI | 42 | 30 | | M | | | |
| Cyclone Microsystems | 24 | 15 | | Macrolink, Inc. | 119 | 64 | 1037 |
| Cypress Semiconductor | CV4 | — | | Maxim Integrated Products | 143 | 81 | |
| D | | | | Mercury Computer Systems | 13 | 8 | |
| Diversified Technology | 46-47 | 33 | | Micro-Sys | 147 | 180 | |
| Dynatam, Inc. | 41 | 29 | | Micron Technology | 55 | 39 | |
| E | | | | Microware Systems Inc. | 35 | 24 | |
| E-Systems, Inc. | 148 | — | | †Military & Aerospace Electronics | 36 | 41 | |
| Electronic Solutions | 1 | 2 | | †Military & Aerospace Electronics | 49 | 84 | |
| Embedded Computer Conference | 141 | 80 | | | | | |
| Engineerium | 146 | 169 | | | | | |
| Eyring | 54 | 38 | | | | | |

ADVERTISERS INDEX

| ADVERTISER | PAGE NO. | CIRCLE NO. | **IDA NO. | ADVERTISER | PAGE NO. | CIRCLE NO. | **IDA NO. |
|----------------------------|----------|------------|-----------|------------------------------|----------|------------|-----------|
| †Military & Aerospace | | | | T | | | |
| Electronics | 86 | 58 | | *Texas Instruments..... | 33 | — | |
| Mitsubishi Electronics/ | | | | Texas Microsystems, Inc..... | 80 | 52 | |
| EDG | 16-17 | 10 | | Themis Computer..... | 79 | 51 | |
| Mitsubishi Electronics/ | | | | Toronto MicroElectronics .. | 122 | 68 | |
| EDG | 28-29 | 18 | | | | | |
| Mitsubishi Electronics/ | | | | V | | | |
| EDG | 40 | 28 | | VenturCom | 34 | 23 | |
| Mitsubishi Electronics/ | | | | Vigra, Inc. | 82 | 53 | |
| EDG | 67 | 46 | | Vista Controls | 117 | 63 | 1052 |
| Mitsubishi Electronics/ | | | | VME Microsystems Int'l .. | 115 | 62 | |
| EDG | 69 | 47 | | Vail Silicon Tools, Inc..... | 146 | 166 | |
| MIZAR | 56 | 40 | | | | | |
| Motorola TSD Group..... | 51 | 36 | | W | | | |
| N | | | | Brian R. White Company...27 | 17 | | |
| NEC Electronics, Inc..... | 109 | 90 | | Wescon '92 | 139 | 79 | |
| Needham's Electronics..... | 147 | 175 | | Wind River Systems..... | 65 | 45 | |
| Neuro Dynamx..... | 27 | 16 | | | | | |
| Nohau Corporation | 146 | 167 | | Z | | | |
| O | | | | Z-World Engineering..... | 146 | 172 | |
| Oasys, Inc. | 45 | 32 | | Ziatech Corporation..... | 83 | 54 | |
| Octagon Systems | | | | Zwick Systems..... | 147 | 183 | |
| Corporation | 147 | 177 | | | | | |
| OEM Integration..... | 149 | 83 | | | | | |
| Omnibyte Corporation..... | 128A | 119, 120 | | | | | |
| OrCAD Systems..... | 147 | 179 | | | | | |
| P | | | | | | | |
| Pentek | 84 | 55 | | | | | |
| Phase Three Logic | 148 | 185 | | | | | |
| Q | | | | | | | |
| Quantum Software Systems, | | | | | | | |
| Ltd. | 43 | 31 | | | | | |
| R | | | | | | | |
| Radstone Technology | 39 | 27 | | | | | |
| Ready Systems | 60 | 42 | | | | | |
| Robotrol Corporation | 147 | 182 | | | | | |
| S | | | | | | | |
| SBE, Inc..... | CV3 | 88 | | | | | |
| Sophia Systems & | | | | | | | |
| Technology | 77 | — | | | | | |
| Spectrum Signal | | | | | | | |
| Processing..... | 127 | 70 | | | | | |
| Synergy Microsystems..... | 2 | 3 | | | | | |
| Synergy Microsystems..... | 21 | 13 | | | | | |

**Computer Design's
INSTANT DATA ACCESS
ADVERTISERS**

To receive data sheets instantly from these advertisers via FAX, dial (617) 494-8338 and when prompted, enter the 4 digit document number shown here.

| ADVERTISER | Page NO. | IDA NO. |
|---------------------|----------|---------|
| Brooktree..... | 75 | 1026 |
| Hitachi..... | 86 | 1034 |
| Macrolink | 119 | 1037 |
| Vista Controls..... | 117 | 1052 |

* Domestic only
† International only
** These advertisers have made arrangements to provide you with instant product data via Computer Design's Instant Data Access (IDA) System.

The Advertisers Index is published as a service. The publisher does not assume liability for errors or omissions.

* Domestic only

† International only

** These advertisers have made arrangements to provide you with instant product data via Computer Design's Instant Data Access (IDA) System.

The Advertisers Index is published as a service. The publisher does not assume liability for errors or omissions.

2ND ANNUAL ANALOG & MIXED-SIGNAL DESIGN CONFERENCE.

HYATT REGENCY
San Francisco Airport
OCTOBER 28-30, 1992
BURLINGAME
CALIFORNIA

100% DEDICATED TO YOUR NEEDS

The Analog & Mixed-Signal Design Conference is dedicated to the

unique needs of

mixed-signal design. If you

are an analog design engineer, a digital designer

who must address the analog implications of high speed design, or an engineering manager,

you should attend this conference.

You'll learn about analog and mixed-signal simulators, different designs techniques, analog implications of high-speed systems, and much more such as:

- How to bring mixed-signal design to market quickly
- New innovative solutions to tough design problems
- The capabilities and limitations of analog and mixed analog/digital ICs, ASICs and CAE/CAD tools
- Secrets to boosting your productivity from the best in the business
- Cost effective ways to implement high-speed digital, mixed-signal, and analog designs

JOIN YOUR COLLEAGUES

Don't miss this opportunity to strengthen your skills and increase your designing power.

For more information call or fax Tom Loftus at the numbers below.

Tom Loftus

ANALOG & MIXED SIGNAL

☎

415-905-2354

FAX

415-905-2220

DESIGN CONFERENCE

When you need to reach design and development qualified engineers and engineering managers in the electronics OEM, think *COMPUTER DESIGN* first.

Only COMPUTER DESIGN:

•serves today's \$300 billion computer based electronics OEM

•provides design directions, options and choices—with exclusive "why-to" editorial

•delivers 100,000 engineering managers and engineers—100% design and development qualified!

•reaches over 70,000 design and development engineering managers—more than any other design publication in the market

•is the fastest growing publication in the field

For more information on how to subscribe or advertise, please contact: Tim Tobeck, National Sales Manager/Associate Publisher at (508)392-2116

COMPUTER DESIGN *Technology and Design Directions*

OneTechnology Park Drive
Westford, MA 01886
(508) 692-0700

CIRCLE NO. 87

COMPUTER DESIGN

Technology and Design Directions

Executive Office: One Technology Park Dr, PO Box 990, Westford, MA 01886
(508) 692-0700 Fax: (508) 692-7780 Telex: 883436

Publisher: David L. Allen (508) 392-2111

Associate Publisher/Editor-in-Chief: John C. Miklosz (508) 392-2114

Associate Publisher/National Sales Manager: Tim L. Tobeck (508) 392-2116

Administrative Assistant: Peg Alexander (508) 392-2112

Marketing Specialist: Betsy Anderson (508) 392-2209

Ad Traffic Manager: Kelly Rice (508) 392-2198

Recruitment Advertising, Postcards/System Showcase: Sue Shorrock (508) 392-2185

Circulation Director: Bob Dromgoole (918) 832-9213

Circulation Manager: Paul Westervelt (918) 832-9287

List Rental: Deanna Rebro (918) 831-9551

Reprints: June Bozarth (918) 835-3161

SALES OFFICES

Associate Publisher
Tim L. Tobeck

National Accounts Managers

Eastern/Midwestern Region

Tim Pritchard
One Technology Park Dr
PO Box 990
Westford, MA 01886
Tel: (508) 392-2217
Fax: (508) 692-7780
Telex: 883436

Southwestern/Northwestern Region

Eric Jeter
19627 I-45 N, Suite 110
Spring, TX 77388-6030
Tel: (713) 353-0309
Fax: (713) 288-8350

Northern California

Diane Palermo
910 Campisi Way
Suite 1B
Campbell, CA, 95008
Tel: (408) 371-7551
Fax: (408) 371-6329

Strategic Accounts Representatives

Eastern U.S.

Sue Nawoichik
One Technology Park Dr
PO Box 990
Westford, MA 01886
Tel: (508) 392-2118
Fax: (508) 692-7780

Western U.S.

David Singer
One Technology Park Dr
PO Box 990
Westford, MA 01886
Tel: (508) 392-2109
Fax: (508) 692-7780

U.K./Scandinavia

David Round
Westmead House
123 Westmead Road
Sutton
Surrey SM1 4JH
Tel: (081) 770-1100
Fax: (081) 770-9779

France/Belgium/Southern Switzerland/Spain/Netherlands

Daniel R. Bernard
10, rue Michelet - B.P. 279
78502 Sartrouville Cedex, France
Tel: 01-39-14-67-80
Fax: 01-39-14-70-14

Germany/Austria/Northern Switzerland/Eastern Europe

Johann Bylek
Verlagsbuero Johann Bylek
Stockaekerring 63
D-8011 Kirchheim/Muenchen
Federal Republic of Germany
Tel: 89-903-88-06
Fax: 89-904-35-26

Italy

Luigi Rancati
Rancati Advertising
Milano San Felice Torre 5
20090 Segrate, Italy
Tel: 02-70300088
Telex: 328601 RANCAD I
Fax: 39-270300074

Japan

Toshio Egusa
PubliNetwork
C407, 2-22-6
Tsukuda Chuo-ku
Tokyo 104, Japan
Tel: 03-3536-5404
Fax: 03-3536-5490

Southeast Asia

Taiwan/Hong Kong/Korea/China/Thailand/Philippines/India/Pakistan/Singapore/Malaysia/Indonesia/Australia/New Zealand
Tom Gorman
CCI Asia-Pacific, Ltd
Suite 905, Guardian House
32 Oi Kwan Rd, Happy Valley
Hong Kong
Tel: 833-2181
Fax: 834-5620

COMPUTER DESIGN (ISSN-0010-4566) is published monthly by PennWell Publishing Company, 1421 S Sheridan, Tulsa, OK 74112. Second-class postage paid at Tulsa, OK 74112 and additional mailing offices. Editorial offices are located at One Technology Park Dr, PO Box 990, Westford, MA 01886. Subscription Prices: Free to design and development qualified engineers and engineering managers in the U.S. and Canada. Paid to all others. Qualified engineers and engineering managers outside the U.S. and Canada - air shipped - \$85 one year. For non-qualified recipients in the U.S. - \$88 one year; in all other countries - air shipped - \$154 one year. Call (918) 832-9263 for subscription information. Microfilm copies of COMPUTER DESIGN may be purchased from University Microfilms, a Xerox Company, 300 N Zeeb Rd, Ann Arbor, MI 48106. POSTMASTER: Send change of address form to COMPUTER DESIGN, Circulation Department, PO Box 3466, Tulsa, OK 74101. ©1992 COMPUTER DESIGN by PennWell Publishing Company. All rights reserved. No material may be reprinted without permission from the publisher. Officers of PennWell Publishing Company: Frank T. Lauinger, Chairman and Chief Executive; Joseph A. Wolking, President and Chief Executive Officer; John Ford, Senior Vice-President; Carl J. Lawrence, Senior Vice-President; V. John Maney, Vice-President/Finance; W. Steve Zimmerman, Vice-President/Corporate Services.

CLASSIC VME

Reliability and Performance from SBE



Turn your product into a classic with SBE's complete selection of intelligent, high-performance VME communications controller boards. All offer the unrivaled quality and reliability you expect from SBE, a world-leader in VME...a company committed to taking VME to the highest level with unsurpassed technical support, on-time delivery, and leading-edge product development.

SBE enjoys a reputation for innovation in all major communications technologies—High-Speed Serial, Token Ring, FDDI, and Ethernet.

SBE's fully featured VCOM-34 with 4 or 8 ports of high-speed serial commu-

nications offers speeds to T1 and E1, with optional support for X.25, frame relay, and SS7 communications protocols.

Whether you manufacture mini/supercomputers, workstations, or high-performance data communications products, count on SBE to deliver premium features, without a premium cost. And this includes integrated hardware/software solutions, comprehensive development assistance, and on-going support throughout a product's life cycle.

For classic VME performance and reliability to meet your most demanding LAN and WAN interface requirements, there's only one place to turn: SBE.

For fast action call 1-800-347-COMM

Germany: (+49) 8061-3240

SBE, Inc., 2400 Bisso Lane, Concord CA 94520
See us at Interop Booth #2138

SBE
Communications &
Real-time Solutions

FAST FLASH 22V10



EXTRA! CYPRESS STUNS WORLD WITH FIRST FLASH PLD.

Stop the presses! Once again, Cypress has the lead story in PLD technology for high-performance systems. Cypress is first on the world scene with 10 ns, Flash 22V10 devices. Electrically alterable 22V10s are your fastest route to risk-free inventory and ease of design. Cypress scoops the competition again!

Also newsworthy: This 22V10 is CMOS, needing just 90 mA max (commercial) and 100 mA (military applications), so it stays cool for reliable operation. Choose from DIP, PLCC and LCC packaging options.

Cypress's Flash 22V10 is the latest member in a complete family of landmark PLD products with the widest variety of speeds, densities and architectures to suit your application. Read all about it—call the Cypress hotline for your free Flash sample certificate and data sheet today.

FREE FLASH SAMPLE HOTLINE: 1-800-858-1810*
Ask for dept C117.

*In Europe, fax your request to the above dept. at (32) 2-652-1504 or call (32) 2-652-0270. In Asia, fax to the above dept. at 1 (415) 961-4201.
© 1992 Cypress Semiconductor, 3901 North First Street, San Jose, CA 95134. Phone 1 (408) 943-2600, Telex: 821032 CYPRESS SNU UD, TWX: 910-997-0753.



PAL C 22V10 D



MAYBE YOU JUST NEED BETTER DRIVERS.

Stuck in a SCSI data-transfer bottleneck?

A first-class SCSI-2 VME board *could* get you out of that jam...*if* getting that board up and running doesn't leave your whole engineering team in gridlock.

That's why, when you buy a Rimfire SCSI-2 VME board from Ciprico, you don't just get a product. You get a partner.

As your partner, we'll help you write the drivers you need, and provide on-site installation. We'll even work with your engineering team to customize hardware or software for your special application needs.

And you get Rimfire's unmatched performance, quality and features. Like the new Rimfire 3870 and 3590 – the first VME boards with independent two-channel SCSI-2 architecture.

First-class hardware *and* first-class support and service. *That's* how you beat bottlenecks. And *that's* what Ciprico is all about.

TURN TO CIPRICO WHEN YOU'RE IN A SCSI JAM.

CALL 1-800-SCSI-NOW


Ciprico, Inc.
2800 Campus Drive
Minneapolis, MN 55441
Fax: (612) 551-4002

CALL (44) 635 873666

Ciprico (Europe) Ltd.
7 Clerewater Place, Lower Way
Newbury, Berkshire RG13 4RF England
Fax: (44) 635 871996

CIRCLE NO. 1

 CIPRICO



Tiburon's Greg Pope on:
Software Testphobia

COMPUTER DESIGN

*Technology
and Design
Directions*

\$5 USA
\$6 Canadian
®

FOR ELECTRONIC ENGINEERS & ENGINEERING MANAGERS

Analog designers still trail behind their digital counterparts

**Performance analysis spots
hardware/software bottlenecks**

**Windows demands drive
PC chip set features**

**High-speed A-D converters
shift to new architectures**



Special Report on Future Computing

**Is neural computing the key
to artificial intelligence?**

